Journal of Intelligent Manufacturing (1996) 7, 39–53

Expert system support for environmental assessment of manufacturing products and facilities

STEPHEN D. KLEBAN,[1] GEORGE F. LUGER[2] and RANDALL D. WATKINS[1]

[1]*Sandia National Laboratories, Albuquerque, NM 87185, USA*
[2]*Department of Computer Science, University of New Mexico, Albuquerque, NM 87131, USA*

Received October 1994 and accepted April 1995

The goal of environmentally conscious design for manufacturing is to select materials and processes that minimize environmental impact. This paper describes a general and uniform way to analyze the environmental impact of manufacturing based on the product decomposition, the materials used in the manufacturing processes, and the particular view of the environment. To accomplish this task, we developed a computer program, called EcoSys^{TM}, that assists manufacturing engineers and environmental reviewers in assessing the environmental consequences of their manufacturing decisions.

Keywords: Design for environment, environmentally conscious manufacturing, expert systems, knowledge-based systems, life cycle analysis

1. Introduction

The goal of our project is to identify and quantify material consumption and waste generation within the manufacturing process: that is, to measure the general environmental impact of the manufacturing of products. Because of the complexity of considering the multiple components and processes of a created product in the context of environmental issues, maufacturing process engineers are in need of a tool to assist in their decision making. EcoSys^{TM} provides the ability to evaluate the environmental impact of manufacturing, at both the product and process level. This environmental analysis is made in the context of multiple environmental models.

For example, manufacturing processes within a facility can be evaluated and ranked according to environmental consequences. Another example is that an entire product can be analyzed to determine which of its parts or processes most affect the environment. The system requires facility-specific data for the parts and processes that make up the manufacturing task. Associated with each process is a measure of its material input and waste streams. These product-specific quantities are required for accurate analysis of the manufactured product. The user may also want to design his or her own manufacturing process and compare its environmental impact with that of a currently existing process. To accomplish this, the user need only specify the constituent materials and quantities and evaluate this newly defined process in comparison with others in the context of the environmental models.

In Section 2 we describe the sets of environmental constraints supporting our analysis. Because of the often subjective nature of these constraints we discuss their creation and use in some detail. In Section 3, we present our methodology for the analysis of manufacturing products and facilities. In Section 4, we describe the knowledge representation and constraint propagation schemes in EcoSys. Section 5 offers a complete overview of the EcoSys program and its components. Section 6 presents a number of example analyses of products and processes, including the comparison of individual processes across several environmental models.

2. Environmental life cycle decision analysis models

Within the manufacturing industry, *environmental quality* is a loosely and subjectively defined term (Daly and Cobb, 1989; Piasecki and Asmus, 1990; Greve and Smith, 1992). To conduct meaningful environmental assessments, one must have clearly defined and accepted metrics. In addition, environmental models must determine how environmental consequences can be reported in an objective sense, other than by only the magnitude of some use activity (consumption or emission levels). Another important issue is to define a methodology that compares the impact from very different manufacturing activities.

For instance, we can measure 'tons' of CO_2 emissions from a coal-fired plant, or 'acres' of rainforest consumed. These numbers can be viewed as the 'impact' of a particular process. The actual impact, however, is the measure of degradation of the entire ecosystem as a result of the presence of additional CO_2 or absence of trees. Furthermore, one might expect that the environmental consequences from CO_2 emissions and rainforest destruction are numerous. The impact of either one might be directly comparable to that of the other: that is, each affects global CO_2 balance. Alternatively, the impacts of CO_2 and rainforest destruction might be very different: that is, the result might be the forced extinction of some species in the forest. Unfortunately, defining 'impact' and assigning judgments between dissimilar environmental models has historically led to unmanageable public debate, resolved only in philosophical arenas with little or no influence on current environmental policies. This issue *must* be addressed if we are to have success in promoting effective environmental life cycle work.

A challenge for the environmental management community is to identify and quantify the relationships of material and energy consumption, and waste generation and management, with environmental quality. This systematic process, referred to as *valuation* (De Neufville, 1990), can be conducted using one of a number of techniques including monetization, direct comparison, and abstract comparison. The choice of method will influence to a great extent the framework for the analysis methodology. Only recently has the debate begun over which technique(s) is(are) most appropriate for the valuation of life cycle inventories (Field *et al.*, 1994).

We employ in our environmental modeling the principles of *life cycle analysis* (LCA) as defined by SETAC, the Society of Environmental Toxicology and Chemistry (Fava *et al.*, 1991), and EPA, the Environmental Protection Agency (Vigon *et al.*, 1992). Our approach is structured around what is termed the *product life cycle inventory*. It is important to note, however, that the analysis methodology presented in this paper does deviate substantially from these more traditional life cycle inventory guidelines. Although inventory data generated from the manufacturing phase of the product life cycle are explicitly considered in our analysis (manufacturing is the phase where the greatest degree of data precision can be maintained), full life cycle inventories are considered only in a more qualitative sense. EcoSys has a new approach to LCA in that it assigns environmental attributes to materials that relate to the traditional LCA. Figure 1 shows how the material attributes relate to the life cycle stages.

Traditionally, life cycle inventory has a very specific meaning. The inventory identifies and quantifies the inputs to industrial systems, such as resources and energy, and the outputs, such as air emissions, waterborne effluents, solid waste and other environmental releases incurred throughout the life cycle of a product or process (Fava *et al.*, 1991). Unfortunately, gathering and characterizing life cycle inventory data is fraught with problems, including the tremendous breadth of direct and indirect material and waste streams associated with a single process or product, as well as the associated difficulties of determining appropriate system boundaries. Life cycle inventory analysis has been criticized because of the large economic burdens of 'accurate' assessments. In response, the LCA community has developed strategies to cut corners, such as promoting industry-averaged data, and in doing so has now to deal with significant concerns over data quality. Finally, the emergence of agile, or flexible, manufacturing has complicated the issues of performing inventory analysis. More often the processes once used to produce a single product are now being used to produce a wide range of products and have more sporadic duty cycles. Determination of life cycle inventories for these aggregates of products requires a methodology that accommodates the 'many to many' relationships among products, processes, and materials.

The approach described below was adopted in response to these concerns of traditional life cycle inventory analysis, and also, in large part, because an environmental impact assessment methodology was needed that goes well beyond the characterization of input and output streams. As an important further extension of the current state of life cycle analysis, we define *environmental impact analysis* models and have developed a prototype expert system to present environmental information effectively to the design and manufacturing community. This tool is not strictly environmental compliance-driven, but recognizes that problems associated with environmental damage extend well beyond the constraints imposed by federal and local regulations. Beyond compliance, the possibilities for the creation and use of the concept of 'green design' are exciting.

We began our effort to identify the criteria that affect environmental quality at the process and product levels. We first selected a number of designers, manufacturers, environmental safety and health personnel, and environmental technology staff, and assessed their knowledge of environmental criteria. We solicited perspectives from industry, universities, and the EPA. The results of these assessments, obtained through a survey and a panel

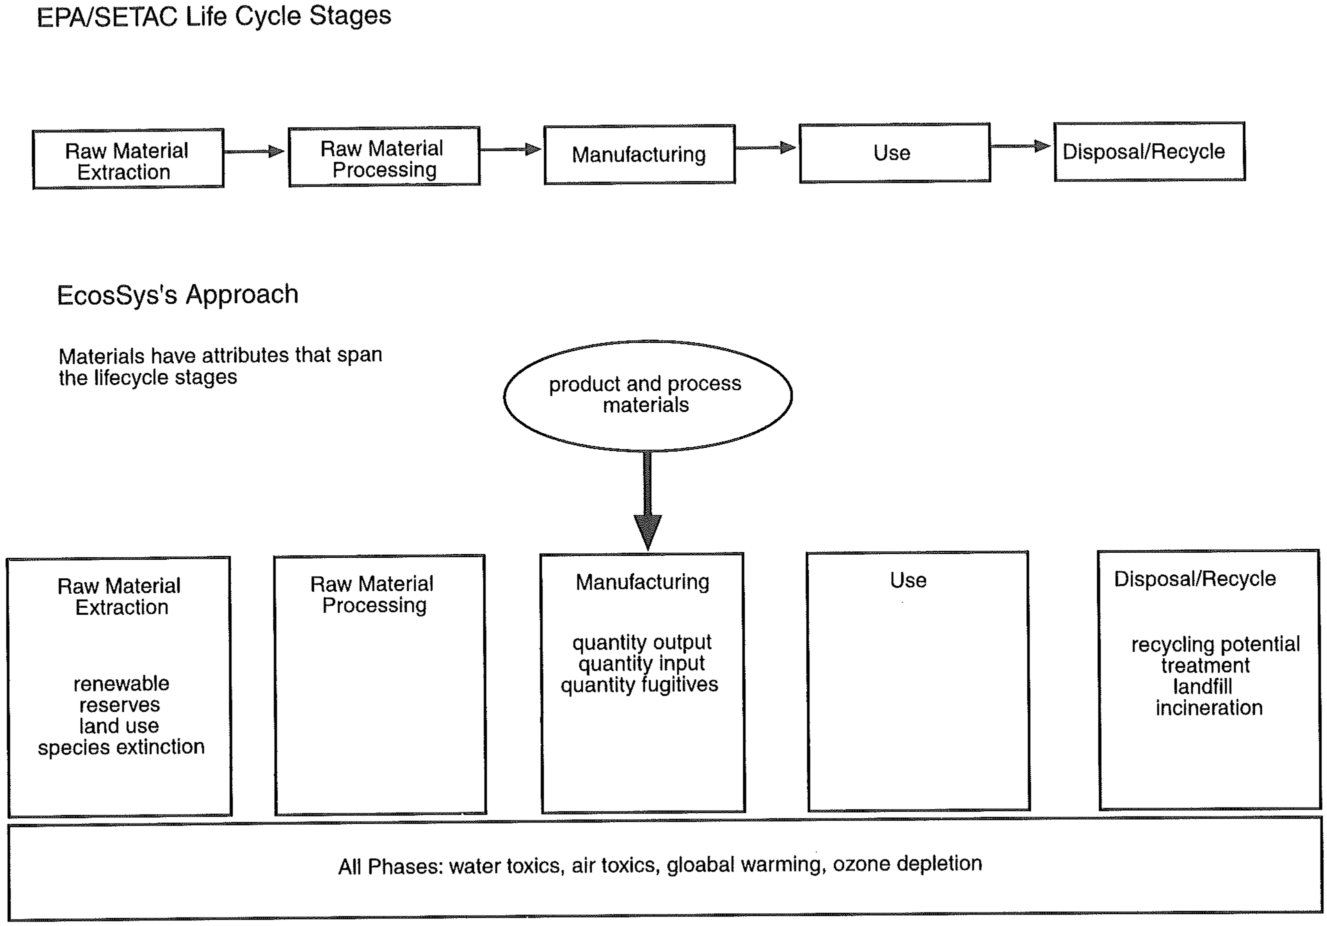

Fig. 1. The relationship of EcoSys material attributes to the EPA/SETAC life cycle stages.

session, were applied in our development of environmental impact decision models that support life cycle analysis. These models, based on the application of the *analytic hierarchy process* (Saaty, 1990; Dyer and Forman, 1992), can be represented in a tree-like structure consisting of goals, objectives, criteria, and alternatives. The basic features of the models are illustrated in Fig. 2. The integration of the environmental models into the expert system is described in more detail in Section 4.

Although there seems to be a general consensus on criteria to include in an environmental impact model, there is significant disagreement among experts regarding the relative importance of the various criteria within and between models. Therefore, we decided to construct a set of models, each supporting significantly different views, and then let the design engineer compare and contrast the advantages and disadvantages of each approach.

These models, based in part on Colby's (1990) different views of environmental management, recognize the subjective nature of environmental quality. Colby proposes five alternative models of the relationship between economic growth and the environment. These define a diverse range of views of environmental quality and ecological health:

(1) *Frontier economics.* This view focuses on economic growth and a free market economy, with little regard to environmental consequences;

(2) *Environmental protection.* Here the environment is viewed as an economic externality that must be safeguarded through laws and regulations. The primary issue in this view is that man simply produces too much waste. This highly anthropocentric view focuses on pollution prevention and waste minimization;

(3) *Resource management.* In this view the environment is again an economic externality, but it must now be internalized in measures of economic performance. The problem according to this model is that man is managing resources poorly. This view recognizes that consumption as well as waste generation affects environmental quality, but is still concerned primarily with effects related to human health;

(4) *Eco-development.* This view stresses the co-evolution of man and ecosystems on an equal basis.

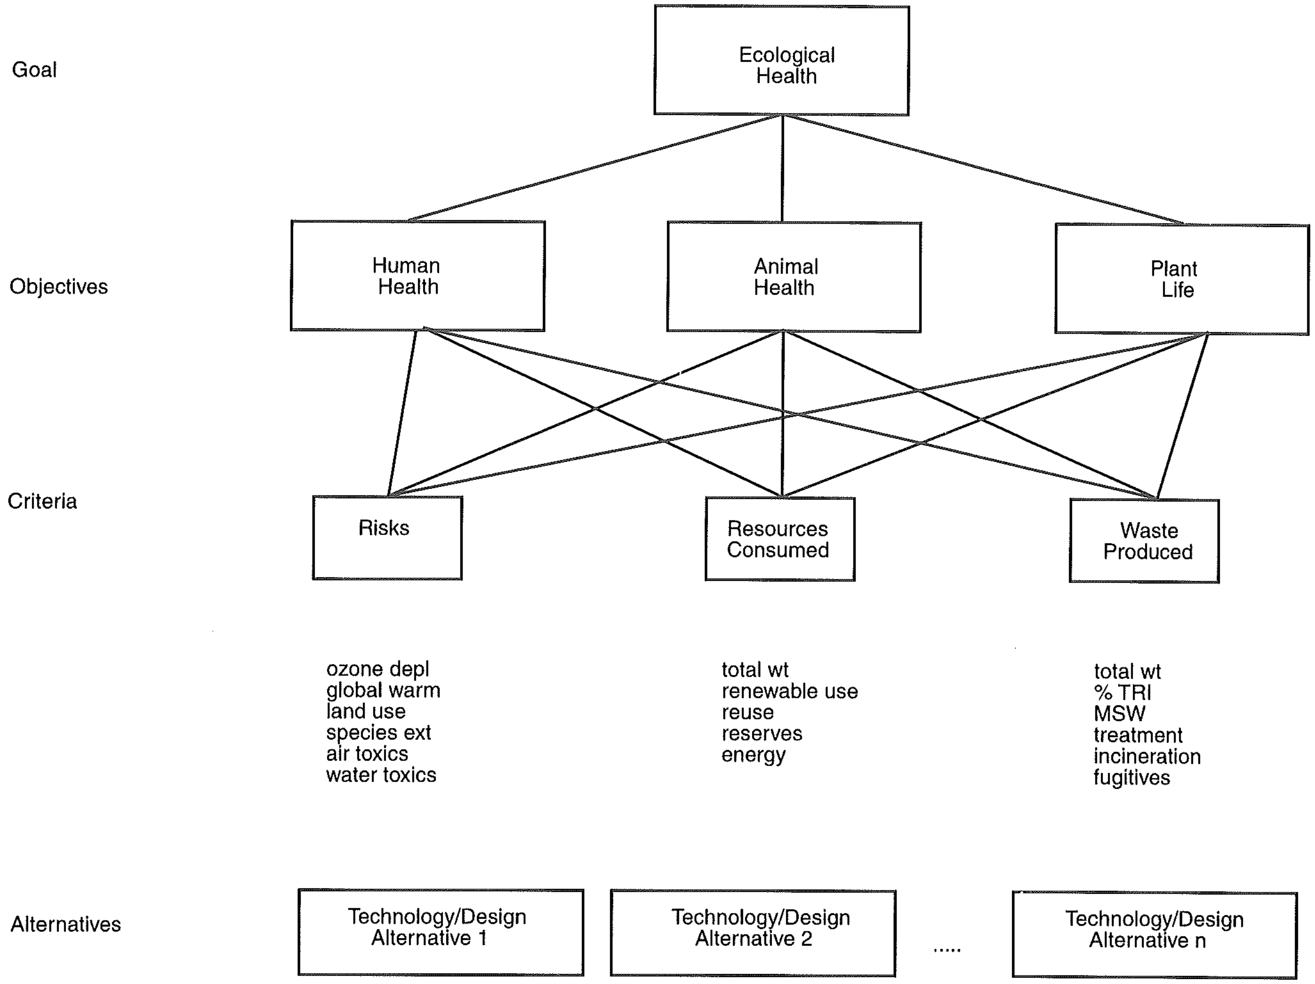

Fig. 2. Environmental model stating goals, objectives, criteria, and alternatives. For example, the ecological health of Alternative 1 can be compared with the ecological health of Alternative 2.

This view suggests that, at present, the scale of economic growth is inconsistent with the long-term coexistence of man with nature;

(5) *Deep ecology.* This view focuses on harmony with nature, and emphasizes drastic reduction in human population and the scale of economic growth.

We feel that the most relevant views, based on the current climate in the environmental policy arena, are *environmental protection, resource management*, and *eco-development. Deep ecology*, although it evolved as a reaction to *frontier economics*, did in fact provide a significant awareness that resulted in the formulation of these other more practical views. Based on this assumption, we generated three decision models consistent with these three intermediate positions. The actual differences between each of the three specific views in the model are represented in how the weights are assigned for the goals and process criteria, as we see in detail in Section 5. A more in-depth explanation of environmental life cycle decision analysis models may be found in Watkins *et al.* (1995). As an example, the environmental protection view considers only human health as its primary objective with no weight given to plants and animals. On the other hand, the eco-development model views human health, animal health, and plant life as equally important.

Criteria selected for inclusion in these models are based on (1) environmental risk, (2) waste production, and (3) resource consumption, with each model viewing the relative importance of these criteria differently. Subcriteria under risk were assigned following the findings reported by an EPA Science Advisory Board study (EPA/SAB, 1990) on risks associated with global environmental problems. Issues including global warming, ozone depletion, land use, and species extinction were identified by the EPA Science Advisory Board as posing the greatest threat to environmental quality. Therefore, these criteria were given greater weight in our models.

The focus of environmental impact analysis is the *process*. The fundamental assumption is that environmental impacts result directly from human activity (processing) rather than from the materials or products themselves. The process defines what materials are required, and in what quantities. The impact from products results from a unique combination of processes or, even more generically, 'events' experienced during the life cycle of that product. Changing the materials used in the product or in the processes will alter the relative environmental impact. Despite this assumption, the basic building blocks for a process are, of course, materials. Materials, although they are viewed as not having direct environmental influences, do have what we call *environmental attributes*. Consistent with the environmental constraint models described above, these attributes of materials include a percent hazardous (toxic release inventory), a disposal option (landfills, incinerate, treat), whether they are a renewable or nonrenewable resource, reuse or recycle, reserves index, energy consumption, and risks associated with ozone depletion, global warming, land use, species extinction, and air and water toxics.

We constructed a materials library to assign values to each of the environmental attributes for a set of selected materials. For *ozone depletion potential* absolute values are applied where available. For *disposal option*, the material is given a designation for the most common (or actual) route of disposition. For all other attributes, we adopted a rating system that assigns a value of 1 if the material has *low* impact for that attribute, 5 for a *moderate* level of impact, and 9 for *high* impact. The rationale for our use of such a discontinuous rating scale is that these values are often derived from limited and sometimes conflicting literature sources, rather than from detailed life cycle studies. As more and more studies are completed, these data will be modified and be made more consistent with a more objective life cycle perspective.

The material library database is used to construct the process definition. From the life cycle inventory, we generate mass balances for material use in and out of processes. To define the process, the appropriate material data sets are retrieved from the library, and are assigned quantities into the process and quantities out of the process (*qtyin* and *qtyout*). There is also a measure of quantities emitted in an uncontrolled manner, termed fugitive (*fugtvs*). Quantities are used to derive (1) the total quantity of materials consumed by the process, (2) the total quantity of waste produced by the process, and (3) the weighted averages of the environmental attributes for the process. To define the product, all of the processes applied are then summed over the life cycle of the product. This summation is conducted in much the same manner as summing the materials that constitute a process.

Application of environmental impact attributes to materials, as just defined, results in an impact analysis methodology that deviates from more traditional life cycle analysis guidelines. Rather than conducting full life cycle inventory assessments of all materials used, we are providing inventories only for manufacturing. The choice of environmental attributes was selected to permit the creation of a qualitative life cycle *perspective* without the need to collect complete life cycle inventory data. Again, while inventory data can be used to help generate the ratings, these data are not explicitly required. Our approach will support the evolution of a more 'graded' analysis methodology, where a high degree of data certainty is maintained within the manufacturing analysis.

3. The analysis of manufacturing processes and facilities

From the inception of the EcoSys project, the program designer's goal was to build a program that assisted the human expert in making supported, environmentally sensitive decisions in product design and selection of manufacturing processes. To understand this complex problem-solving task better, we employed the knowledge extraction process from human experts working in the area (McGraw and Harbison-Briggs, 1989). We interviewed a number of experts in three distinct domains: the environment, design, and manufacturing. During these sessions we asked the experts to talk through their problem-solving processes for a number of typical problems. The audiotaped interview sessions were fairly relaxed, with experts asked to go through typical situations. We also asked several different experts to go through the same situations to see the different ways in which human experts focused on similar problems. The program designers were the interviewers and were initially rather naive about the entire product-realization process. An important component of these interview sessions was for the human expert indirectly to instruct the program designers in the context of describing actual problem solutions.

There were three main results of this effort. First, we found that the extraction of environmental knowledge was particularly difficult because of the subjective nature of the problem. Nonetheless, the system required environmental models to support decisions in the design of product or process. This was accomplished through the interviews with environmental experts. These results were described with the creation of the environmental models of the previous section.

The second part of the solution involved a great deal of interaction with manufacturing engineers and product designers. Traditionally, design and manufacturing are two very distinct steps in the product realization process. We needed to understand this decoupling before we could integrate the two sequential steps into the product/process structure described in Section 4.

The final result of our interviews was to identify the reasoning steps that allowed the human expert to tie all this information together. The first consequence was to gather together the part and process information in a database together with environmental information about materials. We next describe the nature of the reasoning that brought the system together.

Two important insights about the human problem solving in this domain came from these interview sessions. First, we found that the skills of the human expert could be represented rather well in the context of the rule-based solution paradigm. In fact the relationships between product and process information and the constraints of the environmental models translated directly into *if. . . then. . .* rules. We offer examples of several of these in Section 4. The second result of the interviews was the determination that the human experts' problem solving was strongly *goal-directed.* This goal-directed behavior could be seen from the fact that through most of their problem solving the human experts were trying to accomplish some task. Examples are as simple as 'If I want to use their solder compound, what cleaning solvent and flux is required?' or 'If I substitute another subcomponent, can I cause less CO_2 emission?' Indeed, it took very little reflection to convince the programming team that modeling the conjunction of components and processes under environmental constraints required goal-directed rule processing.

This identification of the problem solving as goal-driven played an important role in the design of our search algorithm. In our prototype, this meant the selection of the M.4 (1993) software, as is discussed further in Section 4. A final point is that the rule-based representation for joining processes, materials, and components in the context of environmental constraints translated very conveniently into an *and/or* graph for problem solution search. The goal of our prototype was to design a program that offered a first approximation for the knowledge-intensive and constraint-based approaches of the manufacturing and environmental experts. Once we were satisfied with our prototype (see discussion of Section 4), the second version of our program was built in C++.

We utilized three types of rule to implement the manufacturing and environmental constraints. The first describes goals or problems to be solved: either a manufactured part or a manufacturing process. The structure of these rules was derived from the interaction we had with the manufacturing engineers and designers. The second type of rule controls the information flow to and from the environmental models. The third rule type integrated environmental as well as product and process information from the database for the analysis. This is shown diagrammatically in Fig. 3. The rules that describe the part, process, and material integration will be described in Section 4.

The processing logic that resulted from our analysis turned out to be very close to well-understood and often already implemented processing in manufacturing: that is, the and/or trees are very close to product structure trees used to develop bills of materials and to perform materials planning. The process operations are close to operations files used for capacity planning, shopfloor scheduling, and loading (Bray, 1988). As generic product and process analysis tools exist, they often do not interrelate the product material requirements with the manufacturing process material requirements. In the context of environmental analysis, both material requirements must be considered. Commercial software could have been used to model the product and processes but further development would be needed to combine the information and add the environmental analysis. Although we did not choose this route, since our first goal was proof of

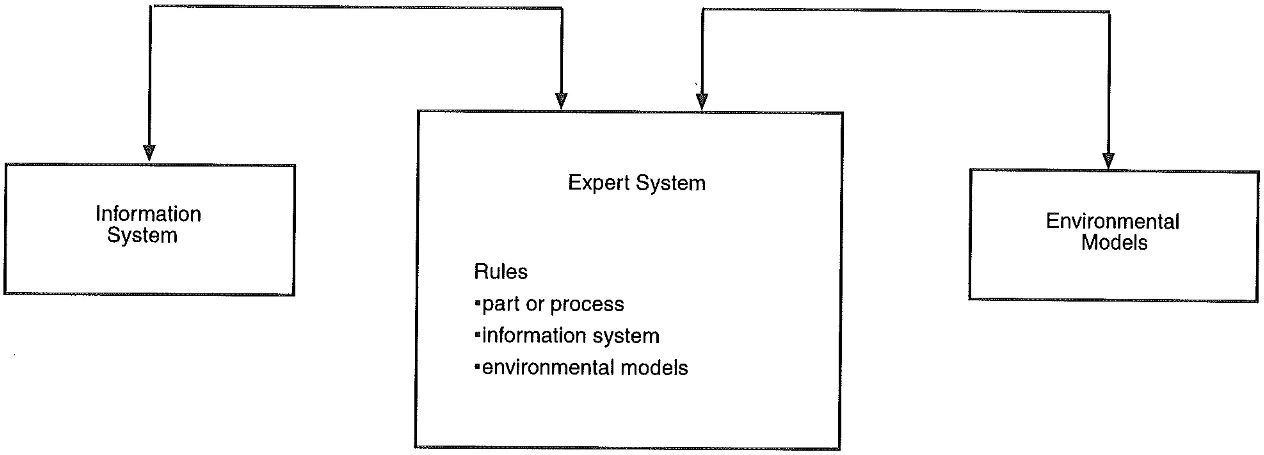

Fig. 3. Expert system rules interacting with the information system and environmental models. The information system and environmental models are external to the expert system: that is, only the expert system rules interact with them.

concept, it is valid and probably the preferred route for industry.

4. Knowledge representation and constraint propagation in EcoSys

4.1. *The and/or tree*

Traditionally, the representation for product fabrication is in the form of some type of flow diagram similar to that of Fig. 4. These diagrams are similar to organizational charts that outline the hierarchy of employees at a workplace. At times, they are difficult to read when long lines connect the boxes with no apparent order for their composition. The hierarchical relationships are sometimes confusing, as boxes may be connected horizontally as well as vertically.

Based on our analysis of the reasoning of human experts, and the natural decomposition of this expertise into *if*. . . *then*. . . rules, we selected an alternative representation to the flow diagram, that of the *and/or* tree. These trees not only clearly define the hierarchical relationships within a product structure but can also represent alternative ways to construct the product. The product structure in Fig. 4 is shown as an *and/or* tree representation in Fig. 5. This is actually an *and* tree because there are no alternative choices in the fabrication process. Everything is *and*ed together to make the final product. (We recommend Luger and Stubblefield (1993) for descriptions of AI data structures, including *and/or* trees and *data-driven* and *goal-driven* models of *state space search*.)

An *and/or* tree can be used to represent the manufacturing processes that produce the assemblies. Processing nodes in the tree combine parts and assemblies to make a higher-level assembly. The process node also includes input materials for the process, such as chemical cleaners or solder. Figure 6 shows an *and/or* tree that includes

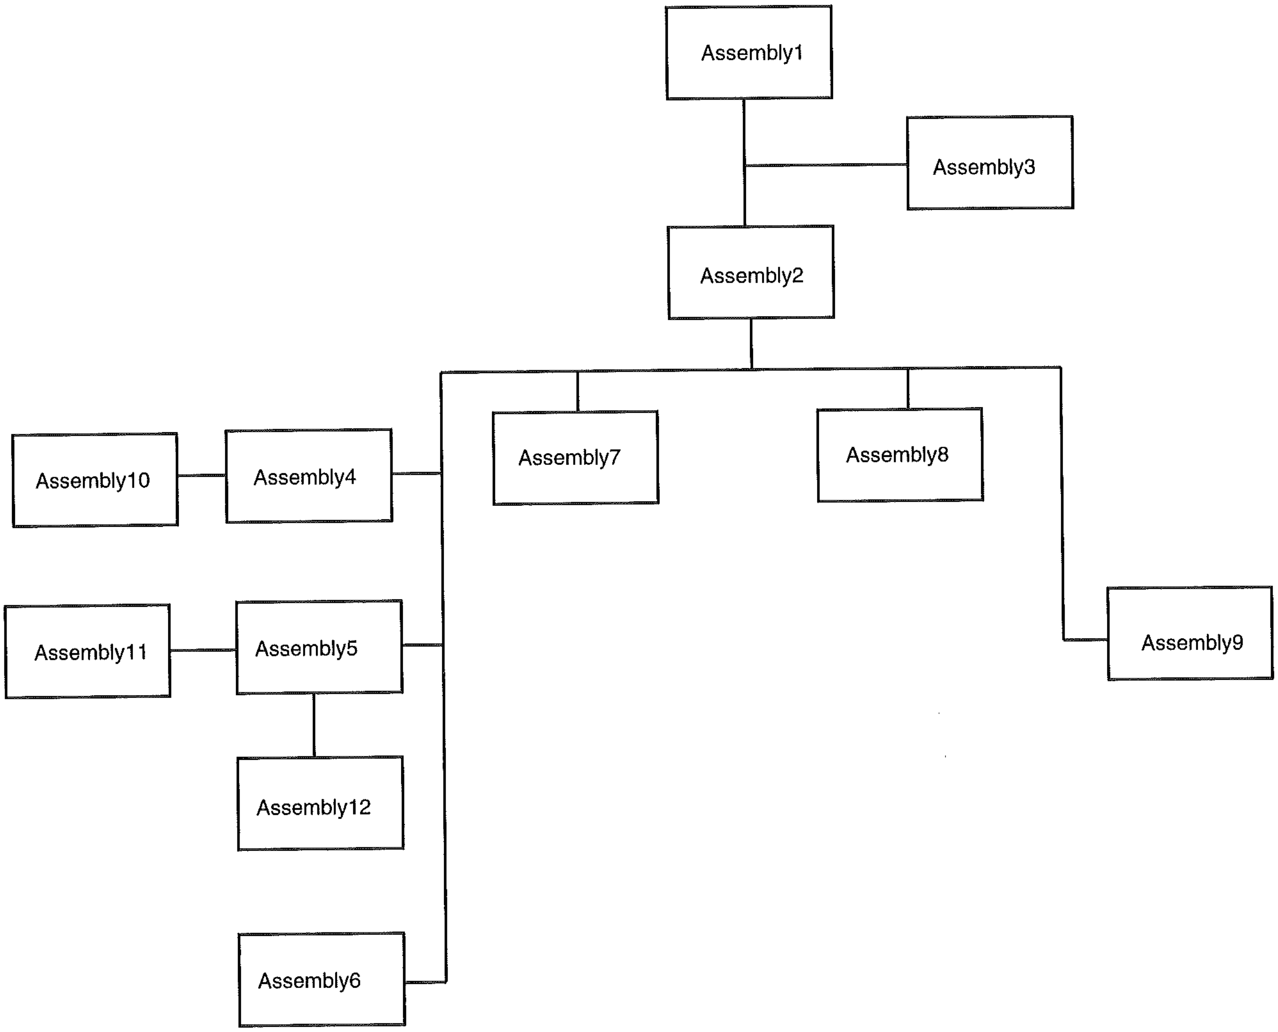

Fig. 4. A typical product flow diagram for manufacturing.

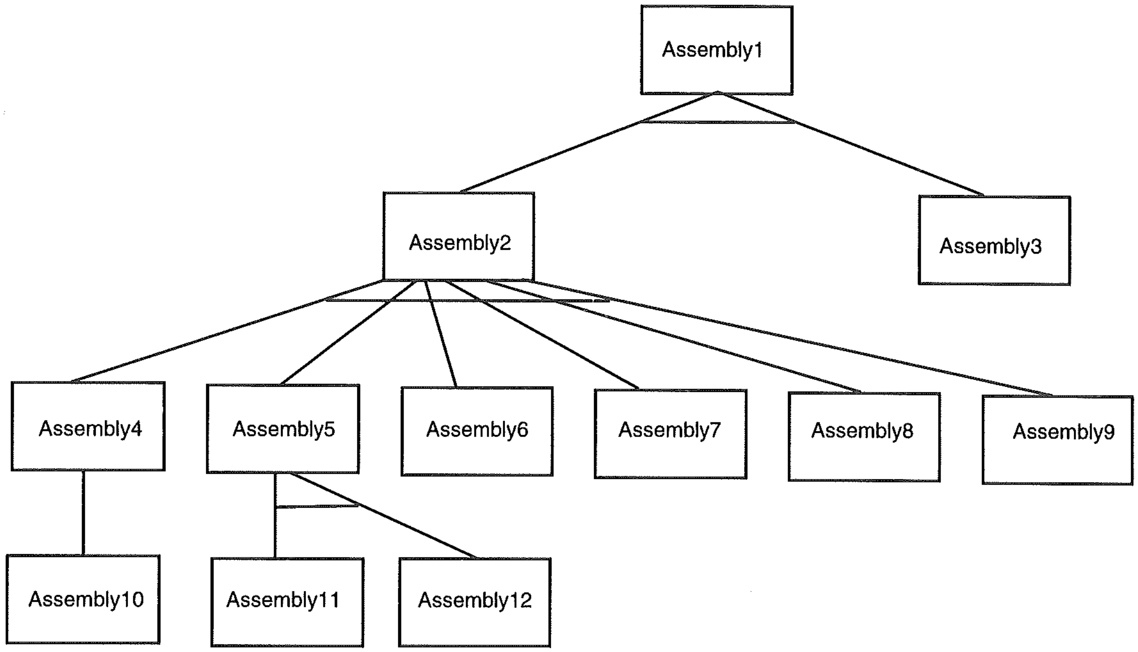

Fig. 5. Product structure of Fig. 4 represented as an *and/or* tree.

process information with product structure. The processing nodes are shown enclosed in rectangles with rounded corners, with materials shown as a single box and assemblies shown as double boxes. This is done for illustrative purposes only. Notice the *or* branch under Assembly 5, where this assembly can be made by either Process6 or Process6a.

Our expert system traverses this tree in a depth-first manner to propagate and combine information used in the analysis of the product. The *or* branches represent alternative parts or processes, and our analysis compares alternatives to determine the optimum combination according to our different environmental models. Handling constraints, such as those of our environmental metrics, requires exhaustive search of subtrees of parts and processes. This meant that a full depth-first or best-first search was required. For computational reasons, such as limited depth of the search, no recursion, and the efficiency of depth-first search in general (Luger and Stubblefield, 1993, Ch. 4), we used depth-first search. Depth-first search is a very reasonable scheme for this problem domain. Our experts would typically perform comparative analysis on small components or on small sets of manufacturing process. They were not interested in large problems like ‘Build me an environmentally friendly computer’. Their reasoning was more ‘Give me analysis of the process involved in making a printed wiring board’.

As part of the early prototyping process, we asked manufacturing and environmental engineers to evaluate our reasoning processes and the results. The human experts were very comfortable with this automated reasoning process, and felt it reflected their own reasoning in fairly complete detail. (We used several of the protocol analysis tools common in the cognitive science literature; see Luger, 1994, Ch. 9.)

4.2. *The analysis of constraints with rule-based software*

The rules in our expert system represent the relationships between the nodes and their children in the *and/or* tree, and generally represent the reasoning of a human expert. For our first system we utilized the M.4 (1993) software product from Cimflex Teknowledge: a strong backward-chaining product for goal-driven search. The rules capture the constraints in the *and/or* tree that reflect the product and processes we wanted to analyze. A typical rule has the format:

IF (A and B and C) THEN D

This is a simple inferential rule that states: if A, B, and C are true then D is true. The rule representation maps directly to the *and/or* tree. For example, the rule that represents the ‘pre-tin’ manufacturing process is

IF (solder and flux and TCE) THEN pre-tin

There is a separate rule for every node in the tree that has subnodes. The leaf nodes of the tree are database entry points for product and process information. This way, the entire *and/or* tree structure that represents the product and process is represented as a set of rules that

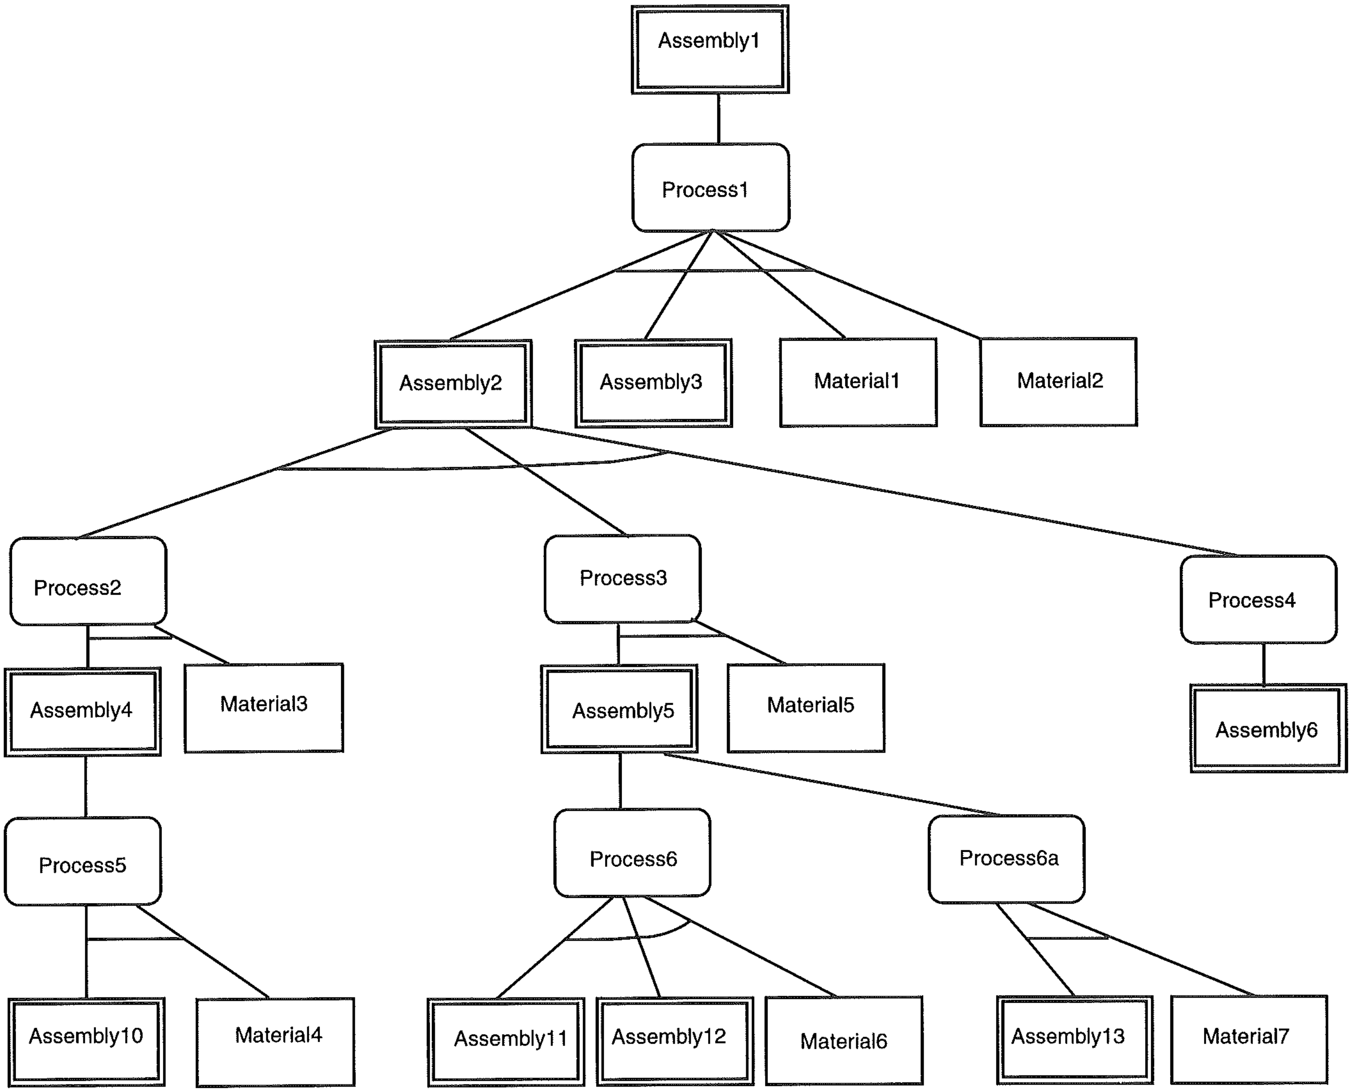

Fig. 6. *And/or* tree showing processes and assemblies; *and* branches are joined by a horizontal line.

ultimately connect the environmental constraints of product and process. Constraint analysis across this data structure is known as a *state space search*.

We used *goal-driven* search of the state space, as we noted in the previous section. Goal-driven search is appropriate for this problem because the types of questions we have with respect to the environmental models of Section 2 are goals to be satisfied. For example, if we want to assess AssemblyA with the *resource management* model, we present the goal to the expert system: 'What are the environmental constraints on manufacturing AssemblyA according to this model?' As the constraint analysis deepens in the tree we find the pre-tin process and present the subgoal: 'What is the environmental cost of the pre-tin process?' Thus the constituents of the pre-tin process are described by the conditions that make up the pre-tin *and/or* rule relationship.

As happens in many expert system projects, once the first version of the program is designed and validated as to its functional adequacy and verified as meeting all the requirements of the end users, the program is rebuilt. This allows the following versions of the program to implement the problem-solving strategies directly and more efficiently, and the sometimes awkward expert system software tool to be replaced. In our situation, the M.4 rule base was replaced by inheritance hierarchies developed in C++.

5. Overview of the EcoSys architecture and subsystems

The architecture of the EcoSys system is presented in Fig. 7. EcoSys runs on a server machine, currently a SPARCstation 10, with users remotely accessing the system from their personal computers, be they PCs, Apple

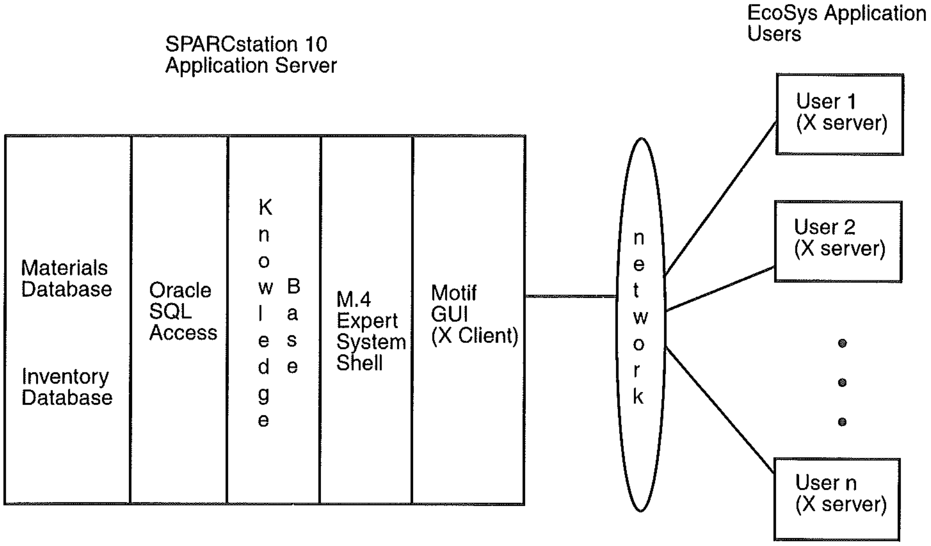

Fig. 7. The EcoSys system architecture.

Macintoshes, or other SPARCstations. The only requirement for the user or client machine is that it have an X-window terminal emulator with access to the network where the server machine resides.

The environmental and inventory information that the expert system manipulates resides in an ORACLE (1991) relational database. It consists of 17 tables, which are joined in various ways to provide usable information for the expert system. There are two main access procedures that the expert system uses. One procedure retrieves environmental impact information, including impact indices for land, water, air, global warming, recyclability, and reusability for all materials. The other procedure provides life cycle inventory information, such as quantities of materials that are used in a manufacturing process.

These access procedures are invoked through the expert system rules. They are simple procedures, which contain Standard Query Language (SQL) statements that manipulate the tables and retrieve the requested data. The *and/or* tree is traversed in a depth-first manner, whereby the information is propagated through the tree to arrive at the desired node for the type of analysis chosen by the user.

The decision analysis module is based on the application of the analytic hierarchy process described in Section 2. The environmental impact model is represented in a tree-like structure consisting of goals, objectives, criteria, and alternatives. The expert system is the heart of EcoSys, encoding the product/process tree of the manufactured product. It has access procedures to the environmental and inventory data and implements the constraints of the environment decision analysis models. It is called from a graphical user interface (GUI), as shown in Fig. 7.

The GUI is written in Motif, a GUI toolkit for the X-window system. It consists of various windows and menus to set control parameters. A portion of the product/process tree for an electronic assembly is shown in Fig. 8. There are four main sections to this user interface screen. A graphical representation of the product structure is depicted with nodes and arcs in the Product Structure section. If the user places the pointer (the mouse) over a node, the information pertaining to that particular node is displayed. The core of the interface is the Tree Browser. The browser allows the user to analyze any node in the product–process tree, not just the entire product. This allows the user to compare and contrast the environmental impact of two or more product subassemblies or manufacturing subprocesses. The Analysis List is simply the list of processes or assemblies to be analyzed, and the Path is the current path in the tree. Once some items exist on the analysis list analyses can be performed, as will be described in Section 6.

From a functional point of view, the user interface controls the rest of the underlying modules: the expert system that manipulates information from the information sources according to the rules defined by the decision analysis module. These results are then reported back to the user and the system waits for more instructions.

6. Samples of EcoSys analysis

EcoSys can model both static and dynamic products and processes: that is, the information pertaining to particular products and processes that currently exist at a manufac-

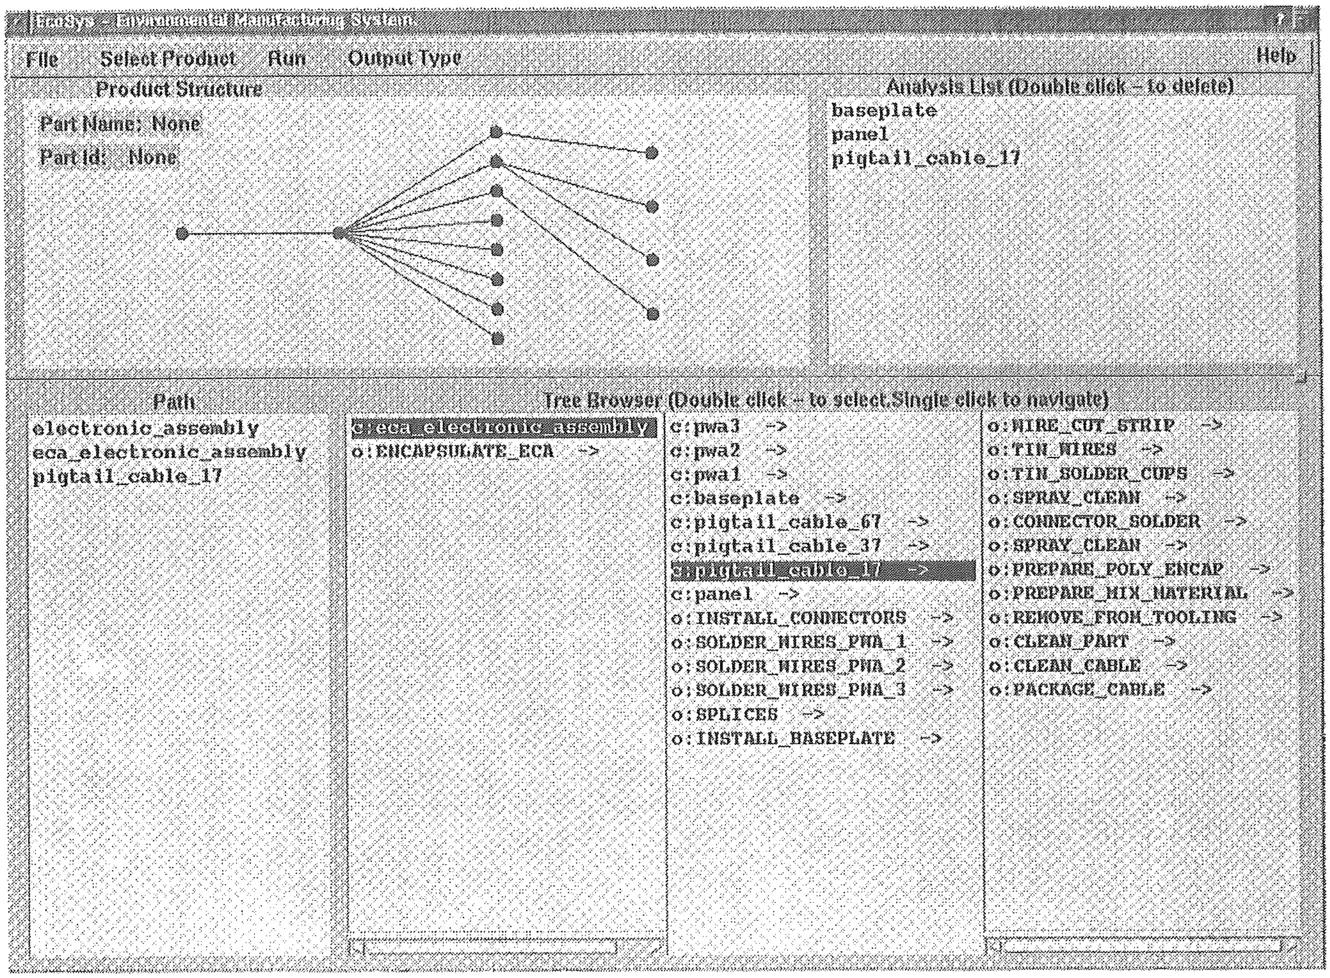

Fig. 8. The main screen of the user interface.

turing facility can be stored statically in a database for retrieval and analysis, or the analyst can dynamically define new products and processes through direct interaction with EcoSys. A process engineer might, for example, want to create a new process for comparison with an existing one. Likewise, a product designer might want to examine the environmental impact of a slight change to an already existing product design. In these cases EcoSys presents the user with a series of data-entry screens to design the new manufacturing process or product. The analysis examples that follow are taken from an existing information database supported by a particular manufacturing facility.

EcoSys has three methods of analysis: material analysis, data summation, and impact analysis. For material analysis, when the user wants to examine information pertaining to specific materials, he or she simply invokes the materials library through the Materials button from the Select Product menu. When this is done, an alphabetical list of all the materials in the database appears in a window of the browser. This window is scrollable, as there are many more materials in the library than can fit in the window. The user scrolls through the window to the desired material and selects it. Any number of materials can be selected for analysis. The user then presses the Run button, and all the environmental information from the

Type	Name	ID	TRI	ISW	Incin	Treat	ODP	ReNu	ReUs	ReSe	Eng	Glob	Land	Spec	Air	Water
material	sodium_car	4553052	0.00	0.00	1.00	0.00	0.00	9.00	9.00	5.00	5.00	5.00	5.00	5.00	1.00	1.00
material	sufuric_ac	4528351	0.93	0.00	0.00	1.00	0.00	5.00	5.00	5.00	5.00	9.00	5.00	5.00	9.00	9.00
material	boric_acid	4528010	1.00	0.00	1.00	0.00	0.00	9.00	5.00	5.00	5.00	1.00	5.00	5.00	5.00	9.00
material	solder	8330203	1.00	1.00	0.00	0.00	0.00	9.00	5.00	5.00	5.00	1.00	5.00	5.00	5.00	9.00

Fig. 9. Examples of specific material information from the database.

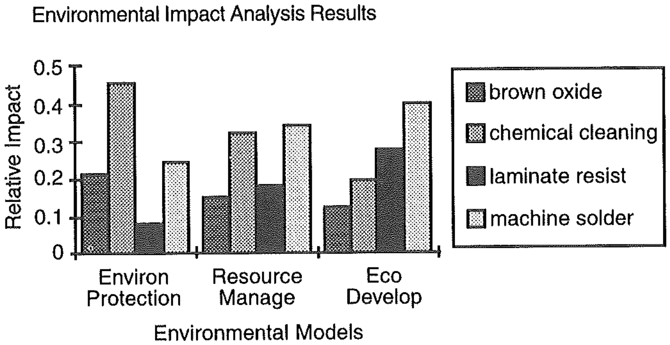

Fig. 14. The results of an impact analysis shown in bar graph form.

from the impact analyses. For example, chemical cleaning yields orders of magnitude more waste than the other processes. Referring back to Fig. 13, the environmental protection view weighs waste production heavily, so this result is not surprising. On the other hand, environmental attributes for chemical cleaning, such as global warming and land use, have low values. These attributes are weighed higher in the eco-development model, contributing to the low relative impact of chemical cleaning according to this view. A facility to explain the results in natural language is planned for the future. There is much detailed data for a curious user to examine, but often the user simply wants the bottom line and an explanation as to how the result was derived.

7. Summary

EcoSys is an automated assistant for manufacturing engineers and environmental assessors who wish to evaluate a manufacturing facility or a manufactured product with respect to the environment. The evaluation methods and algorithms in EcoSys are generic, and to change the application domain only the information system needs to be populated with facility or product data. Thus, to evaluate a new facility, the information system needs to be loaded with the material flow data for the processes at that specific facility. To evaluate a product, material information about the product as well as manufacturing process data must be known. EcoSys currently supports analysis in the electronics industry; we are planning its extension to textiles and chemicals.

Finally, as many of the pieces of component and process analysis are already available in traditional software support packages for manufacturing, the addition of environmental constraints can be seen as a natural extension. In the future, the environmental impact of manufacturing will be seen as but one more standard output for analysis, treated routinely along with the other aspects of cost, inventory, and capacity planning. Integrating environmental analysis from the very beginning of the design process also meets the goals of concurrent engineering (Salzberg and Watkins, 1990).

Acknowledgements

We would like to thank Sandia National Laboratories for support of the EcoSys software development as well as for providing environmental specifications for the products and process of the system. Ron Hall and Larry Claussen, also of SNL, provided the information system design as well as Oracle expertise. Gregory Gershanok of the University of New Mexico provided help in building the graphical user interface in Motif. Bruce Gockel from the AlliedSignal manufacturing facility in Kansas City provided process inventory and product data. Everett Heinonen of New Mexico Engineering Research Institute helped coordinate the EcoSys project. We are grateful for the editorial suggestions from two anonymous reviewers.

References

Bray, O. (1988) *CIM: The Data Management Strategy*, Digital Press, Bedford, MA.

Colby, M. E. (1990) Environmental management in development, World Bank Discussion Paper, World Bank Press, Washington, DC.

Daly, H. and Cobb, J. (1989) *For the Common Good*, Beacon Press, Boston, MA.

De Neufville, R. (1990) *Applied Systems Analysis*, McGraw-Hill, New York.

Dyer, R. F. and Forman, E. H. (1992) Decision support with the analytic hierarchy process. *Decision Support Systems*, **8**, 99–124.

EPA/SAB (1990) *Reducing Risk: Setting Priorities and Strategies for Environmental Protection*, EPA SAB-EC 90-021, Science Advisory Board, US EPA, Washington, DC.

Fava, J., Denison, R., Jones, B., Curran, M., Vigon, B., Selke, F. and Barnum, J. (1991) *A Technical Framework for Life-Cycle Assessments*, Society of Environmental Toxicology and Chemistry, Foundation for Environmental Education, Washington, DC.

Field, F., Isaacs, J. and Clarke, J. (1994) Life cycle analysis of automobiles: a critical review of methodologies. *JOM* (formerly the *Journal of Metals*), **46**(4), 12–16.

Greve, M. and Smith, F. Jr (1992) *Environmental Politics: Public Cost, Private Rewards*, Praeger Publishers, New York.

Luger, G. F. (1994) *Cognitive Science: The Science of Intelligent Systems*, Academic Press, San Diego, CA.

Luger, G. F. and Stubblefield, W. A. (1993) *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*, Benjamin Cummings Publishing, Palo Alto, CA.

McGraw, K. and Harbison-Briggs, K. (1989) *Knowledge Acquisition: Principles and Guidelines*, Prentice-Hall, Englewood Cliffs, NJ.

M.4 (1993) Expert system shell product sold by Cimflex Teknowledge, Palo Alto, CA.

Oracle Corporation (1991) *ORACLE for UNIX Technical Reference Guide*, Version 6.0.30.

Piasecki, B. and Asmus, P. (1990) *In Search of Environmental Excellence*, Simon and Schuster, New York.

Saaty, T. L. (1990) *Multicriteria Decision Making: The Analytic Hierarchy Process*, Vol. 1, McGraw-Hill, New York.

Salzberg, S. and Watkins, M. (1990) Managing information for concurrent engineering: challenges and barriers. *Research in Engineering Design*, **2**(1), 35–52.

Vigon, B. W., Tolle, D. A., Cornaby, B. W. and Lathan, H. C. (1992) *Life-Cycle Assessment: Inventory Guidelines and Principles*, US EPA, EPA/600/R-92/245, February.

Watkins, R., Kleban, S. and Luger, G. (1995) Expert system support for environmentally conscious manufacturing. *International Journal of Environmentally Conscious Design and Manufacturing*, **4**(2), 53–66.