(*AI EDAM*) (1988) **2**(3), 135–168

TREATMENT SELECTION BY CONSTRAINT PROPAGATION A CASE STUDY IN CUTTING FLUID SELECTION

JAMES E. MOGUSH, DOMINIQUE CARREGA, PETER SPIRTES AND MARK S. FOX
Intelligent Systems Laboratory, Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, *U.S.A.*

The GREASE project is an investigation of the application of artificial intelligence to cutting fluid selection and blending for metal machining operations. The problem is to first diagnose the machining operations to determine what fluid characteristics are required, then to select a cutting fluid which satisfies the required characteristics. The problem is exacerbated by the need to select a single fluid to be used by multiple types of operations on a variety of materials. Diagnosis is relatively simple, but treatment specification is difficult due to the variety of operations to be handled.

GREASE uses heuristic search in which the evaluation function is *heuristically constructed*. The construction of the evaluation function begins with the determination of the characteristics of an optimal fluid based on deep knowledge of the machining operations and materials. This is then altered heuristically according to problems diagnosed with the current fluid. Once the evaluation function is complete, it is used to select an existing fluid from the product line. GREASE has been tested extensively with results which equal that of the experts and has been field tested by the Chevron Corporation.

1. Introduction

In 1984, our research group was presented with the problem of designing a system for the diagnosis and treatment of problems related to the use of cutting fluids in the machining of metals. Cutting fluids are used to provide lubrication and cooling, and to prevent the welding of the metal being machined to the machining tool. The problem has two parts: first, to diagnose what is wrong with the current fluid, and second, to select a new fluid which eliminates those problems. Both the diagnosis and treatment are based upon knowledge of the materials to be machined, the machining operations, and known problems with current fluid. At first glance, it appeared that the problem was an instance of heuristic classification (Clancey, 1984), and would be amenable to rule-based diagnosis techniques pioneered by MYCIN (Shortliffe, 1976) and Prospector (Duda *et al.*, 1978). The techniques utilized in these systems have reached such a stage of development that a variety of systems are now in 'production use' in domains such as medical diagnosis (Aikins *et al.*, 1983), turbine diagnosis (Osborne, 1986; Fox *et al.*, 1983), and telephone line diagnosis (Vesonder, 1983).

After further analysis, our initial view turned out to be incorrect. The first part, diagnosing the current fluid to identify solution requirements, is quite simple and amenable to shallow causal reasoning. It is the second part, identifying a cutting fluid which best satisfies these requirements, which requires 'deeper' reasoning. Though the set of fluids from which to select a treatment is relatively small (less than 100) and easily enumerated, it is the combinations of evidence (i.e., the cross product of materials, operations and current fluid problems) and their causal links to the fluids which cannot be enumerated. Nor is it the case that most of the combinations are pathological; actual machining situations can come in almost any combination.

In order to solve this problem, we had to develop a better understanding of *why* fluids correct machining problems. In particular, the approach we have taken to treatment selection involves:

the deepening of the causal representation from the fluid property level (e.g. process→property→fluid) to the representation of fluid chemistries (e.g. process→property→chemistry→property→fluid).

the combination of qualitative and quantitative causal knowledge to represent relationships between fluid properties and their chemistries.

constraint propagation to identify the changes in fluid properties.

the use of a simple form of search where existing fluids are compared to a heuristically constructed ideal fluid.

The result of our approach is a system we called

Received 24 August 1988.

0890–0604/88/030135 + 34$02.00/0

GREASE.[1] GREASE is a heterogenous diagnosis system, in that it combines both heuristic classification with constraint propagation and heuristic selection in designing optimal fluids, and then, selecting the best, matches from a pre-defined set of fluids. It is composed of the following phases:

DIAGNOSIS—interpretation of machining diagnostics and special machining requirements in terms of cutting fluid property inadequacies;

CONSTRAINT GENERATION—specification of compositional alterations to a cutting fluid to treat inadequate property values identified by the DIAGNOSIS phase;

DESIGN—formulation of an 'optimal' cutting fluid to best satisfy the operations on specific materials in a machine shop. The 'optimal' fluid design includes alterations from the CONSTRAINT GENERATION phase;

SELECTION—selection of cutting fluids from a product line which best match the 'optimal' fluid for the shop as determined during the DESIGN phase;

EXPLANATION—interpretation of the reasoning at each phase to provide sound justification for the cutting fluid selections.

The approach taken in GREASE provides two benefits:

1. A straightforward semantic representation of the domain, including cutting fluid chemistries, their properties, materials, and machining operations;
2. Another approach to heuristic selection where an ideal solution is constructed heuristically and then compared to existing candidates. A side effect of this approach is the reduction of the work required to add new fluids to the knowledge-base to a simple database entry; property information is provided without having to create any new causal links.

The rest of this paper provides a detailed description of GREASE. In Section 2 we begin by reviewing the domain. Section 3 provides an overview to the problem solving architecture. Section 4 provides a detailed description of GREASE's knowledge representation. Section 5 steps through the diagnosis and selection algorithms. Results of a set of experiments in selecting cutting fluids relative to an expert and a salesman experienced in cutting fluid selections are described in Section 6. Section 7 compares our approach to those found in other diagnostic systems. Our conclusions are in Section 8.

[1] Could be construed to be an acronym for General Reasoning Engine And Selection Environment.

2. Machining problems and diagnostic issues

In order to understand how cutting fluids are selecte and designed, it is necessary to understand th different functions of cutting fluids, and th characteristics of the cutting fluids, materials bein machined, and machining operations.

2.1 CUTTING FLUID ROLES

Primarily, a cutting fluid contributes in three wa to the machining process. It acts as a lubricant, coolant and an anti-weld agent.

As a lubricant, a cutting fluid reduces the he generated during the machining process by reducin the friction between the workpiece and the cuttin tool. As an antiweld agent, it counteracts the ten ency of the work material to weld to the tool und the heat and pressure generated in the cutting operatio

To perform satisfactorily as a lubricant, the cuttir fluid must maintain a strong protective film at t portion of the area between the tool face and t metal being cut where hydrodynamic conditions ca exist. Such a film assists the chips in sliding readi over the tool. Besides reducing heat, prop lubrication reduces the wear of the tool and lowers t power requirements.

If a cutting fluid performs its lubricating functic satisfactorily, the problem of heat generation from t cutting tool, workpiece and chip is minimized b cooling still remains an important function. T perform this function effectively, a cutting fluid shoul possess a high thermal conductivity. Water has a hig thermal conductivity and is a very effective coolan but its lubricating properties are practically nil. As result, water-based cutting fluids—emulsions—a good coolants but poor lubricants. On the other han straight oils have relatively low thermal conductivit so that they must depend on fluidity for effecti cooling ability; hence, the faster they flow over a operation, the more heat they can absorb and car off per unit of time.

In some instances, extreme temperatures an pressures at the cutting interface cause the chip, segments of it, to weld to the tool face. The build-u resulting from such welding may occur to such degree that the effective tool shape is drastical changed and all phases of the operation are serious affected. To overcome welding, effective antiwel characteristics may be imparted to the cutting fluids b incorporating various additives. These are usual materials, such as fatty oil, sulfur or chlorine, whic by chemical reaction form a surface film of low shear strength at the chip–tool interface. The effectiveness of these chemical films is understood to be limited by their respective melting points.

The primary functions of a cutting fluid appear to be closely interrelated. A cutting fluid which is a good lubricant will generally be a poor coolant, and vice versa. Properly selecting a cutting fluid consists of satisfying these particular requirements for specific machining processes on materials.

A cutting fluid must also satisfy various secondary requirements, less directly related to the machining process, but nevertheless important. A cutting fluid should, for instance, flush chips away from the work area; protect the finished work surfaces, the tool and, the machine against corrosion and stain; should not smoke nor fog in use, nor produce dermatitis; have a pleasant odor; and be ecologically safe and non-toxic.

Special-purpose requirements may also be imposed to a cutting fluid. A 'grinding fluid' involved in a lapping operation, for instance, must act not only as a lubricant, but also as a medium for suspending the abrasive powder.

2.2 CUTTING FLUID PROPERTIES

We distinguish two categories of cutting fluids: cutting oils and emulsions. It is convenient to think of a cutting fluid as the application of one or more 'products'—straight oils or soluble oils—to a machining process. This distinction allows us to define a cutting oil as a straight oil, or a blend of straight oils, and an emulsion as a water-based solution of a soluble oil. An emulsion is thus characterized by both a soluble oil and a dilution ratio.

The effectiveness of a cutting oil is determined by its physical properties which, in turn, are determined by its chemical composition:

The viscosity of a cutting oil affects its cooling and lubricity properties. The greater the viscosity of the oil, the better its lubricating power, and the poorer its cooling performance. Severe machining operations require high viscosity fluids to enable the oils to adhere better to the tool and workpiece. Less severe operations are generally run at higher speeds which create more heat and consequently, utilize lower viscosity fluids since cooling is the most important factor. In addition to better satisfying lubricity requirements, viscous fluids carry metallic chips more easily and help flush them away from work areas.

Another important factor which affects the lubricating power of a cutting oil is its fatty oil percentage. The higher the fatty oil percentage, the greater the lubricating power of the oil.

The total sulfur, active sulfur,[2] chlorine and phosphorus percentages of an oil account for its antiweld properties. By chemical reaction, these additives form a surface film of low shear strength at the chip–tool interface. The effectiveness of these chemical films is understood to be limited by their melting points: iron chlorides are effective up to 600°C; iron sulfides, 1000°C. Cutting oils containing active sulfur are classified as 'active'. They stain copper and its alloys, and cannot be recommended for the machining of such materials.

Secondary properties of cutting fluids include antimist, anti-foam anti-rust, anti-wear, rust inhibitor, corrosion inhibitor and odor masking capabilities.

The dilution ratio associated with an emulstion is a very important factor. It is directly related to the cooling and lubricating powers of the fluid. The greater the dilution ratio of an emulsion, the better its cooling power, the poorer its lubricating performance.

The total sulfur, active sulfur, and chlorine percentages of an emulsion account for its antiweld properties, as in cutting oils.

Unique properties of emulsions include stability of the emulsion and degradability.

2.3 MACHINING OPERATIONS

Most machining operations may be described and understood as variations of a cutting tool shearing a workpiece or material. Metal ahead of the cutting edge of the tool is compressed, and removed from the workpiece in the form of a chip, by a process of plastic deformation and shearing. Chips fall into three basic categories: discontinuous, continuous, and continuous with build-up-edge. In the latter case, a fragment of work material—the build-up edge (BUE)—sticks to the tool in the region of the cutting edge and protects it against excessive wear. However, too large a built-up edge may result in poor surface finish.

There are several hundred machining operations which are variations or combinations of drilling, milling, planing and shaping, turning, and grinding.

Each machining operation, in turn, has several properties characterizing it and which affect cutting fluid selection. These include:

speed of machine tool;
feed rate of material machined;
depth-of-cut;

[2] Active sulfur chemically reacts with copper at a temperature of 150°C.

tool material composition (affecting hardness and brittleness);

geometry and characteristics of cutting operation.

The most common operations can be ranked into classes based upon increasing severity of the operation—this is a rough classification based upon experience and is a complex function of the properties of the operation. Specific machines may not be easily classified into the severity hierarchy since their geometry or operating conditions may be substantially different than those classified. In addition, some machines performed a multiplicity of operations of different severity.

The characteristics that a cutting fluid must satisfy are different for different severity operations as previously mentioned. As a result, selection of a cutting fluid for a multiple-operation machine must satisfy more characteristics.

2.4 MATERIAL MACHINED

The ease that a material can be machined is referred to as its machinability.

The machinability is a function of the machine operation conditions and the material composition and conditions.

The material properties affecting machinability include:

MICROSTRUCTURE—this is the grain structure of a material. Materials with similar microstructures machine similarly. Uniformly of microstructure favors long machine-tool life,

GRAIN SIZE—small grain size renders a metal ductile and easily machined, but makes it hard to obtain good surface finish. Intermediate grain sizes are best.

HARDNESS—the hardness is the material's resistance to indentation. A higher hardness generally results in good surface finish, but usually is less easily machined,

METALLURGICAL CONDITION OF THE METAL DUE TO HEAT TREATMENT—heating and cooling operations change the physical properties, such as the hardness and the microstructure, of the material. These operations include annealing, normalizing, tempering, quenching, etc.,

METALLURGICAL CONDITION OF THE METAL DUE TO WORKING—metalworking operations affect the physical properties of the material. These operations include casting, forging, hot- and cold-rolling, etc.,

COMPOSITION—the chemical composition of the material greatly affects the overall machinability. Different percentages of different elements withi steels, for example, greatly affect all physic: properties such as hardness, ductility, and tensil strength. An AISI classification system relates th material's composition and its approximate machi ability under standard treatment conditions using standard machining operation.

Materials are generally machined in a soft, easil machinable condition. However, due to varyin control in production, variability does exist in th machinability for specific materials.

2.5 FUNCTIONAL REQUIREMENTS

To ascertain the functionality required b GREASE, several meetings were held with sales an research engineers of the sponsor. The followin summarizes the results.

Selection for a particular specification of machinin operation and material. Selection based upon thes specifications corresponds to rare, but ideal an well-understood situations, where much empiric: knowledge has been gathered and compiled int tables. With the use of such tables, the recommend: tion of a satisfactory cutting oil does not require deep understanding of the phenomena involved in th machining process. This type of situation is generall handled by sales engineers. However, a fine understanding of the process integrating various othe parameters may lead to the recommendation of a optimal product, which could differ from the mor general ones recommended by the tables. In othe cases, when operating conditions or other constrain are atypical, a deeper understanding of the phenom ena is required, and an expert must be consulted.

Selection for a range of machining operations an range of materials. This situation represents mor than 70% of selection cases in the field. A machin shop may have a variety of operations to perform o different machines with different materials. Rathe than using a specific cutting fluid for each (operatior material) combination, the best selection of two c three fluids for the shop to result in satisfactor overall performance is desired. Quite often, th selection process consists in satisfying the require ments of fewer combinations involving the mos salient constraints imposed by the sets of operation and materials.

Selection which most closely matches a competitor' product currently used. This procedure requires match of the cutting fluid property for the range c machining operations and materials rather than th specific chemistry of the fluid itself, althoug approximating the chemistry would result in similar properties. This case represents less than 10% of all selection cases encountered in the field.

Selection based upon diagnosis. In this instance, an improved recommendation is formulated based upon the properties of the current fluid utilized and the diagnosis of why the current fluid selection is unsatisfactory. Diagnostics are encountered in about 5% of cases in the field. However, due to salesman psychology, improvements in the current cutting fluid properties are generally promised in order to encourage a cutting fluid purchase.

The functionality required by GREASE is quite diverse. Of course, one could focus on one part of the problem and ignore the rest, the result being a simpler architecture. The approach taken in the GREASE project is to explore the issues surrounding the design and construction of a system which can span all the functions described above. In particular, how are knowledge representation and problem-solving affected when input can take many forms (e.g., properties, signs, symptoms), and the analysis must transform from selection to diagnosis and treatment when an earlier attempt is found to be unsatisfactory.

3. System architecture

The multiple functions of GREASE, including selection, diagnosis, treatment, and explanation are implemented as a series of successive processing stages (Figure 3). GREASE first characterizes the cutting fluid selection problem, performs diagnosis of any known inadequacies of the current fluid, then begins treatment by calculation of an 'optimal' fluid or fluids. Treatment continues as GREASE then evaluates candidate fluids and determines the best fluids which approximate the 'optimal' condition.

Each processing stage in GREASE is summarized as follows:

3.1 SHOP DEFINITION

All data characterizing a cutting fluid recommendation problem are defined during this stage. These include:

a specification of all machine operations for which a cutting fluid is being recommended;
the materials being machined;
the machining processes relating operations and materials, along with their corresponding:
exceptional operating conditions;
machining diagnostics observed;
specification of the currently used fluid and its chemistry, if known;
any special user requirements.

3.2 DIAGNOSIS

This stage diagnoses any problems observed by interpreting them in terms of imbalances in the cutting fluid property values, which are ultimately reflected in the cutting fluid chemistry. A set of constraints on the properties of the fluid to be selected are generated, for example, 'no sulphur', 'no corrosion', 'increase tool life'.

3.3 EVALUATION FUNCTION GENERATION

An evaluation function, which will be used to rate the available fluids, is constructed at this stage. The evaluation function is a combination of two types of constraints.

1. Predicate constraints which must be satisfied by the selected fluid. These are derived directly from the 'no sulphur' and 'no corrosion' constraints.
2. Relaxable constraints which are refinements of property constraints such as 'increase tool life' to the chemistry level, e.g. 'increase lubricity'. These constraints quantitatively specify the target values of the chemical properties and their acceptable relaxations.

The following steps are performed in constructing the evaluation function.

1. For each process, i.e., material-operation pair in the shop, an ideal fluid is selected via table lookup. This fluid specifies the ideal chemistry required by the process, and the granularity of change of each chemical component in order to satisfy an abstract constraint such as 'increase tool life'. Each ideal

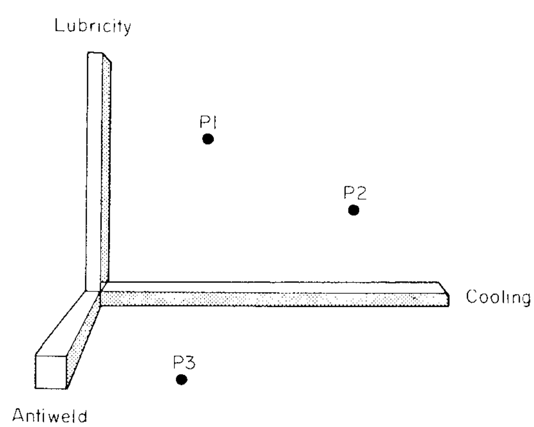

FIGURE 1. Fluid chemistry space

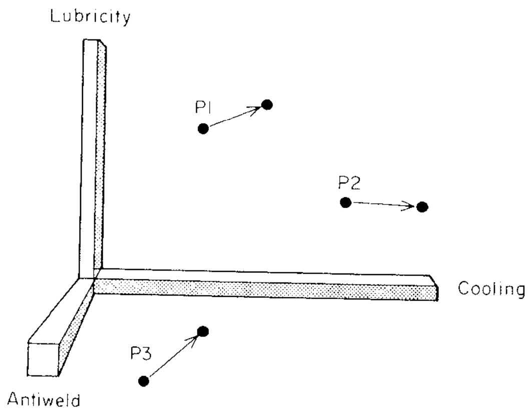

FIGURE 2. Altered fluid chemistry space

fluid represents a point in a space whose dimensions are fluid chemistries: lubricity, cooling, antiweld and viscosity (Figure 1). If a fluid is being used in the shop, that fluid is used as the ideal fluid.

2. Alter the chemistries of the ideal fluids so that they satisfy the constraints generated in the previous stage. This represents a shift of the points in the chemistry space (Figure 2). These are now called the optimal fluids. The knowledge of how to satisfy the constraints, i.e., the amount and direction of change in the chemistry of the fluic was extracted from experts and embedded in th representation of the ideal fluids.

3.4 GENERATION OF CANDIDATES

This stage generates candidate fluids to b evaluated as possible recommendations. In selection this will be the existing product line; if binary blendin is considered, the candidate fluids correspond to a lis of blendable fluids.

3.5 EVALUATION OF CANDIDATES

GREASE then examines each candidate fluid an evaluates it with respect to the 'optimal' fluid for eac machine shop process. The evaluation functio screens the candidate fluids using the predicat constraints, then rates the remaining fluids using th relaxable constraints. The latter evaluation is equiv alent to computing a distance metric from th candidate fluid to the optimal fluids represented in th chemistry space (Figure 2).

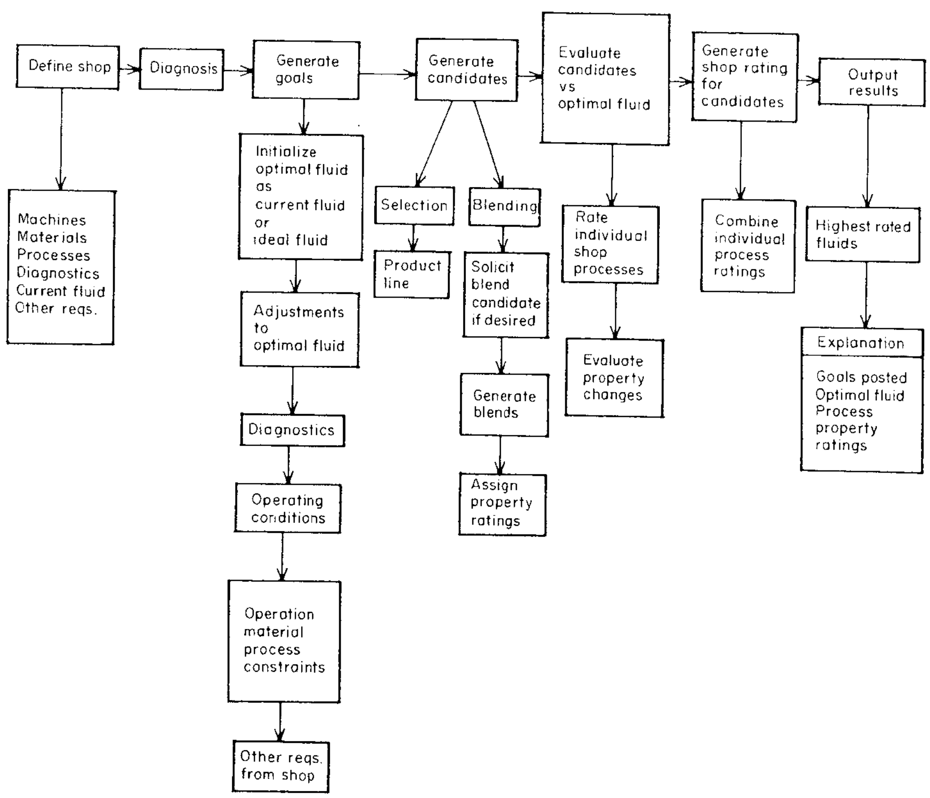

FIGURE 3. GREASE implementation architecture

3.6 GENERATION OF AVERAGE SHOP RATINGS

The shop rating for a candidate fluid is a measure of how well the fluid performs on all machining processes in the shop relative to the other fluids evaluated. The rating is computed as the weighted sum of the ratings for the individual processes defined in the shop.

3.7 OUTPUT RESULTS

The GREASE system finally outputs the fluids in order of decreasing performance, listing how well each fluid performed for each cutting fluid property. Fluids which failed 'fixed-goals' are then listed. An explanation facility provides detailed information on which goals were posted for each machining process, the 'optimal' fluid for each process in terms of its property values, and the rating of the fluid for each process.

GREASE is implemented on a Digital Equipment Corporation VAX computer in Common Lisp and the Knowledge-Craft[3] knowledge engineering system. 'Deep' knowledge of the cutting fluids domain is represented as schemata and the relations among schemata. This provides a flexible and easily comprehensible structure for the implementation of GREASE. It also makes implementation of explanatory capabilities easy. Rather than simply employing an unstructured set of production rules, GREASE uses the 'deep' knowledge of its domain that is embedded in the schemata and their relations to reason from first principles.

4. Knowledge representation

Knowledge is represented as schemata which form a number of taxonomies:

- 'domain' taxonomies representing knowledge about cutting fluids and their application;
- a 'symptom' taxonomy classifying information about possible diagnostics, operating conditions and requirements;
- a 'property' taxonomy, relating high-level cutting fluid properties, such as tool-life and finish, to low level cutting fluid compositional and physical properties;
- a 'goal' taxonomy classifying all possible requirements that can impose on a cutting fluid for a GREASE recommendation.

[3] Trademark Carnegie Group Inc.

4.1 DOMAIN TAXONOMIES

The first stage of GREASE's processing is to extract the shop definition from the user. For each process, the material and operation pair will be used to index into a table of 'ideal' fluids. Acceptable definitions are defined by GREASE's knowledge base of materials, operations and fluids. This knowledge base defines not only what they are, but their characteristics and constraints.[4]

4.1.1 Materials

There are three taxonomies in which a material participates. The first is a simple material definition taxonomy. Materials within GREASE are limited to those which are machined with cutting fluids and are categorized into a ferrous group (i.e. steels) and a non-ferrous group of metals and alloys (Figure 4). Non-metallic substances, such as plastics, etc., are not included. The materials are also considered to be in a 'machinable' annealed condition, except when they are subjected to 'grinding' operations, where they are considered to be in a hardened condition. The schema **STEEL** forms the root of the ferrous materials taxonomy, whose nomenclature is based on the SAE indexing scheme. The taxonomy branches into nine material categories modeled by the schemata **1XXX** to **9XXX**:

```
{{4XXX
  IS-A: STEEL
  USE-OF: MOLYBDENUM
  IS-A + INV: 48XX 46XX 43XX 41XX
  NAME: MOLYBDENUM-STEELS}}
```

Each schema is characterized by two slots: NAME and USE-OF. The NAME slot holds the name of the steel category, and the USE-OF slot records the name of the primary elements occurring in the steel alloy composition. The schema **4XXX**, for instance, denotes the class of 'molybdenum steels', containing the 'molybdenum' element. One level deeper in the STEEL taxonomy, schemata such as **41XX** model subclasses of these steel categories. Finally, individual steels are represented as terminal schemata of the STEEL taxonomy.[5]

```
{{B-111
  IN-GROUP: GROUP1
  IS-A: 11XX
  MACHINABILITY: 94}}
```

[4] Information contained in these taxonomies was derived primarily from Gulf Oil internal documentation (Gulf Oil, 1981*a*, *b*; Gulf R&D, 1982), other publications (American Society of Metals, 1968; Machinability Data Center, 1972) and personal conversations with the Gulf cutting fluid experts.

[5] The current GREASE implementation contains 193 distinct steels.

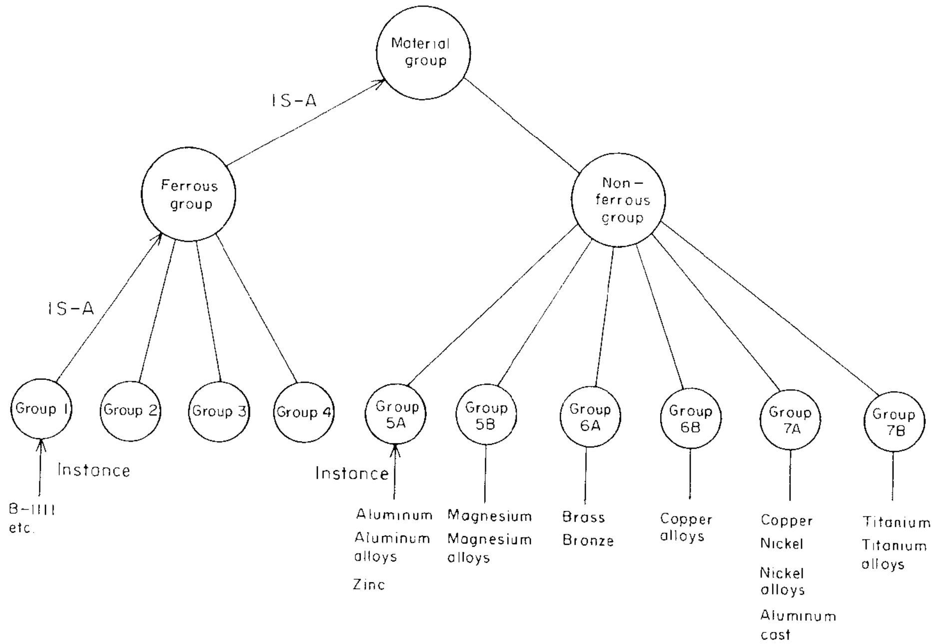

FIGURE 4. GREASE material taxonomy

Each schema is characterized by a MACHINABILITY and an IN-GROUP slot. The first slot holds the machinability rating of the steel, and the second slot records the corresponding machinability group. The individual steels are also the terminal schemata of the MATERIAL-GROUP taxonomy related by the IN-GROUP relation.

The second material taxonomy groups materials which are similar in machining difficulty and require similar cutting fluids. The ferrous materials are divided into four sub-groups (group1 → group4), each of which includes similar steels in terms of machinability. The non-ferrous group is divided into six subgroups (group 5a → group7b), also based upon machinability characteristics and similar cutting fluid composition requirements, such as 'no sulphur' in the fluid.

The machinabilities associated with the material groups are as follows

Group#	Machinability
Ferrous	
GROUP1	Easy—100–70% based on 100% for B1112
GROUP2	Easy/moderate—malleable and cast irons
GROUP3	Moderate—70–50%
GROUP4	Difficult—<50%
Non-ferrous	
GROUP5a,5b	Easy—>100% based on 100% for leaded yellow brass
GROUP6a,6b	Easy/moderate—below 100%
GROUP7a,7b	Difficult—below 100%

GREASE makes no machinability distinction between materials within the same groups, even though they might possess somewhat different machinability values. GREASE would generate similar recommendations for these materials unless the materials generated different material constraints. An example of this would be the generation of a 'no-chlorine' constraint by copper in Group7a, but no corresponding constraint by nickel which is also in Group7a.

The schema MATERIAL-GROUP forms the root of the 'group' classification of the materials taxonomy:

```
{{MATERIAL-GROUP
  USE-OF:
  GROUP-CONTAINS:
  IS-A + INV: NON-FERROUS-GROUP
    FERROUS-GROUP
  ACTIVATES:
  LOW:
  HIGH:}}
```

The slot USE-OF records the names of chemical elements in the composition of the materials. The relation GROUP-CONTAINS, with inverse IN-GROUP, holds the names of the terminal schemata representing the actual materials that correspond to members of the **MATERIAL-GROUP** set. The slots LOW and HIGH define the machinability range characterizing the material group and is represented as numeric values relative to a reference material. The slot ACTIVATES contains goals or 'material constraints' generated by the various materials.

{{**GROUP1**
IS-A: FERROUS-GROUP
MEMBER-OF: CLASS1
GROUP-CONTAINS: B-8735* B-8637* B-8635*
B-8632 B-8630 . . .
HIGH: 100
LOW: 70}}

The third type of material classification is the **MATERIAL CLASS**. They have been defined to group ferrous and non-ferrous materials into sets which have similar machinabilities, but not necessarily similar composition and material requirements. See Mogush *et al.* (1986) for more details.

4.1.2 Operations

There are several hundred possible machining operations requiring cutting fluids, but all are considered variations of the basic operation types included in the 'operation' taxonomy (Figure 5).

The **OPERATION** schema is the root of the operations taxonomy:

{{**OPERATION**
HAS-SEVERITY:
IS-A + INV: GRINDING-OPERATION SAWING
SHAPING PLANING . . .
DEPTH-OF-CUT:
SPEED:
TEMPERATURE:
TOOL-DESIGN:
TOOL-MATERIAL:
FINISH:
HYDRAULIC-REQUIREMENT:
LUBRICITY-REQUIREMENT:
COOLING-REQUIREMENT:
ACTIVATES:}}

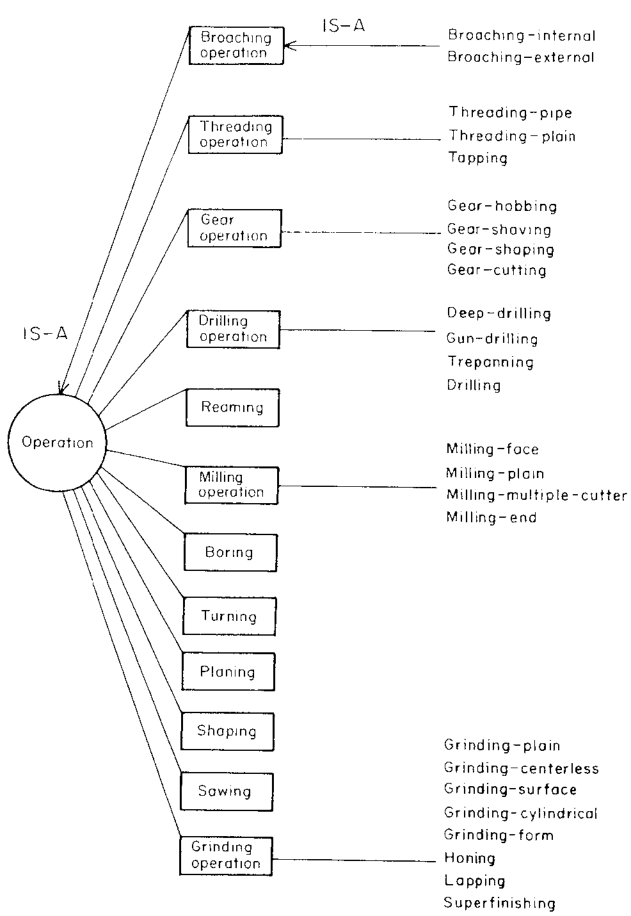

FIGURE 5. GREASE operation taxonomy

The slot HAS-SEVERITY holds the severity class. The slot ACTIVATES records possible constraints that are generated when the operations are present in the machine shop. The remaining slots such as DEPTH-OF-CUT, SPEED, FEED-RATE, COOLING-REQUIREMENTS, etc., record specific properties or machining characteristics of an operation.

Operation types such as **BROACHING-OPERATION**, **THREADING-OPERATION**, **GEAR-OPERATION**, **DRILLING-OPERATION**, etc., are descendants of the **OPERATION** schema and represent the basic operation types in GREASE.

{{**BROACHING-OPERATION**
IS-A: OPERATION
HAS-SEVERITY: SEVERITY1
IS-A + INV: BROACHING-EXTERNAL
BROACHING-INTERNAL
FINISH: HIGH
ACTIVATES:}}

Where operation types are themselves classes, such as **BROACHING-OPERATION**, their descendants, such as **BROACHING-INTERNAL** or **BROACHING-EXTERNAL**, represent the basic operations.

{{**BROACHING-INTERNAL**
IS-A: BROACHING-OPERATION
SPEED: LOW}}

Operations are specified to GREASE corresponding to the specific operation types corresponding to, or 'most similar' to the actual operations desired.

Some operation types in the taxonomy such as **AUTOMATICS-MULTIPLE-SPINDLE** possess more than one parent—in this case **DRILLING-OPERATION**, **TURNING**, **REAMING**, **TAPPING**

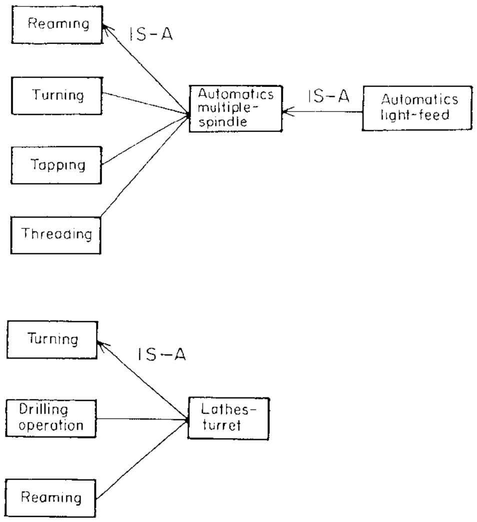

FIGURE 6. Complex operation types

and **THREADING** (Figure 6). Inheritance allows these operations to possess properties of each parent.

The basic operation types fall into ten decreasing severity classes (SEVERITY1 → SEVERITY10) which measures the difficulty of the operation type (Figure 7). Some basic operations which are components of an operation class, such as **DRILLING-OPERATION**, don't possess the same operation severity due to unique characteristics of the particular operation.

4.1.3 Cutting-fluids

GREASE contains two cutting fluid representations:

the Gulf product line fluids which are recommended by GREASE;

'ideal' cutting fluids with cutting fluid property values corresponding to optimal selections for the machining processes embodied in the empirical selection causal network. The property values in the 'ideal' fluids have been tuned by cutting fluid experts.

The cutting fluid product-line taxonomy in GREASE is classified into 'straight-oils' comprising the bulk of the available fluids, 'soluble-oils', and 'chemical-cutting fluids'. This classification is based upon significant cutting fluid behavior differences, and correlations of composition to cutting fluid behavior. Special fluids which are considered 'base oils' and 'blending oils' are not distinguished within the classification, Figure 8.[6]

[6] Specific instances of cutting fluids corresponding to the classification categories are not represented here for proprietary reasons.

Operation	Severity class
Broaching-internal	Severity 1
Broaching-external	
Threading-pipe	Severity 2
Threading-plain	
Tapping	
Gear-shaving	
Gear-hobbing	Severity 3
Gear-shaping	
Gear-cutting	
Deep-drilling	Severity 4
Gun-drilling	
Trepanning	
Automatics-multiple-spindle	Severity 5
Drilling	Severity 6
Reaming	
Milling-face	
Milling-plain	
Milling-multiple-cutter	
Milling-end	
Boring	
Turning	
Lathes-turret	
Automatics-light-feed	
Grinding-form	Severity 7
Planing	Severity 8
Shaping	
Sawing	
Grinding-plain	Severity 9
Grinding-surface	
Grinding-cylindrical	
Grinding-centerless	
Honing	Severity 10
Lapping	
Superfinishing	

FIGURE 7. Operation severity classes

The CUTTING-FLUID schema (Figure 9) is th root of the cutting fluids taxonomy:

The slot TYPE holds the type of the fluid—'insolubl or 'soluble'. In the case of a soluble oil, the sl DILUTION-RATIO records the water ratio. SULFU ACTIVITY flagged as 't' or 'nil' indicates whether th fluid is an 'active' fluid containing active sulfu TYPE-OF-SERVICE refers to the intended duty of th fluid. The slots SUV to FATTY-OIL-PERCENTAGE chara terize the chemical composition and physical chara teristics of the fluid. Additional slots such ANTIMIST-PROPERTY, ANTI-RUST-PROPERTY or ODO MASKANT-PROPERTY refer to specific properti of the cutting fluid. The slots COOLING, LUBRICIT ANTIWELD and ACTIVITY hold measures of cutting flui performance in terms of their fundamental function The slot PRICE holds the relative price of the fluid. Th

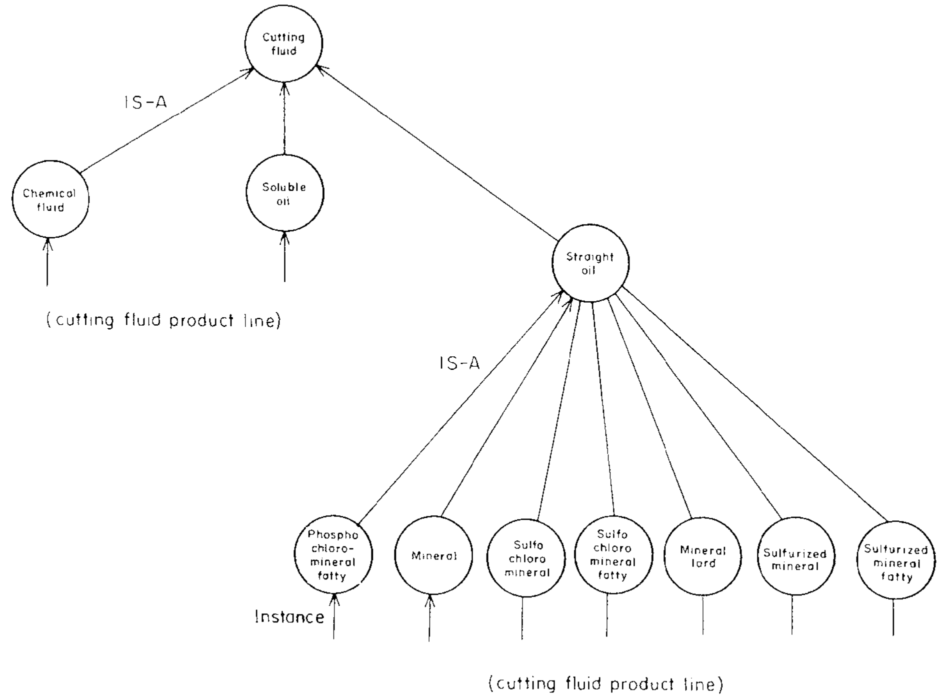

FIGURE 8. GREASE cutting fluid taxonomy

```
{{CUTTING-FLUID
  IS-A + INV: CHEMICAL-CUTTING-FLUID
    SOLUBLE-OIL STRAIGHT-OIL
  TYPE: INSOLUBLE
  DILUTION-RATIO:
  SULFUR-ACITIVITY:
  TYPE-OF-SERVICE:
  SUV:
  KINEMATIC-VISCOSITY:
  TOTAL-SULFUR-PERCENTAGE:
  ACTIVE-SULFUR-PERCENTAGE:
  CHLORINE-PERCENTAGE:
  PHOSPHORUS-PERCENTAGE:
  FATTY-OIL-PERCENTAGE:
  SAPONIFICATION-NUMBER:
  ANTIMIST-PROPERTY:
  ANTI-OXIDANT-PROPERTY:
  ANTI-FOAM-PROPERTY:
  ANTI-RUST-PROPERTY:
  ANTI-WEAR-PROPERTY:
  CORROSION-INHIBITOR-PROPERTY:
  ODOR-MASKANT-PROPERTY:
  DENSITY:
  COOLING:
  LUBRICITY:
  ANTIWELD:
  ACTIVITY:
  PRICE:
  RESULTS:}}
```

FIGURE 9. Cutting fluid properties.

slot RESULTS will record a detailed list-description of the evaluation of the fluid.

Cutting fluids are classified into three distinct types, as previously stated: 'straight-oil', 'soluble-oil' and 'chemical-cutting fluid'.

The **SULFO-CHLORINATED-MINERAL-FATTY-OIL**, **PHOSPHO-CHLORINATED-MINERAL-FATTY-OIL**, **SULFO-CHLORINATED-MINERAL-OIL**, **SULFURIZED-MINERAL-FATTY-OIL**, **SULFURIZED-MINERAL-OIL**, **MINERAL-LARD-OIL** and **MINERAL-OIL** schemata represent seven classes of straight oils. One level deeper, schemata represent the cutting-oils currently available in the Gulf product line. The class of soluble oils represent cutting fluids that are diluted with water in their usage. Specific dilution ratios of these oils are represented as distinct products. The class of chemical cutting fluids is synthetic materials whose chemistry correlates with cutting fluid properties differently than for soluble oils and straight oils.

4.2 SYMPTOM CAUSAL NETWORK

In addition to providing material/operation information, the user can supply symptomatic information. These symptoms are used by GREASE during the diagnostic phase to identify constraints on the cutting fluid. Other conditions that constrain the selection process originate from either the operations or materials specified in the shop definition. These conditions directly identify constraints on the fluid. An example of a material constraint is cast aluminium requiring 'no active sulfur' in the cutting fluid.

{{**ALUMINUM-CAST**
IN-GROUP: GROUP7A
USE-OF: ALUMINUM
ACTIVITIES: (NO-SULFUR 10)}}

There are several symptom types within GREASE:
SHOP REQUIREMENTS are requirements global to all machining processes in the shop such as the 'ecological safety' requirement or the 'no solubles' restriction on the fluids that can be recommended (Figure 10);
PROCESS CHARACTERISTICS are exceptional machine operation conditions such as a 'high speed' operation specific to individual machining processes (Figure 11);
PROCESS DIAGNOSTICS are observed diagnostic conditions for a particular machining process such as 'blued chips'.
Specific symptoms such as **HIGH-SPEED** are instances of the associated symptom type:

{{**HIGH-SPEED**
IS-A: PROCESS-CHARACTERISTIC
ACTIVATES: (TOOL-LIFE COOLING 1.0)
(TOOL-LIFE ANTIWELD 1.0)
STATUS:}}

The slot ACTIVATES indicates which constraints are generated to constrain the selection process. The STATUS slot contains the machining process names that exhibit this symptom.

Process diagnostics are classified into either 'fixed' or 'change' categories. A 'fixed-process-diagnostic' indicates a diagnostic requirement that must be met for the process, such as NO-RUST (Figure 12).

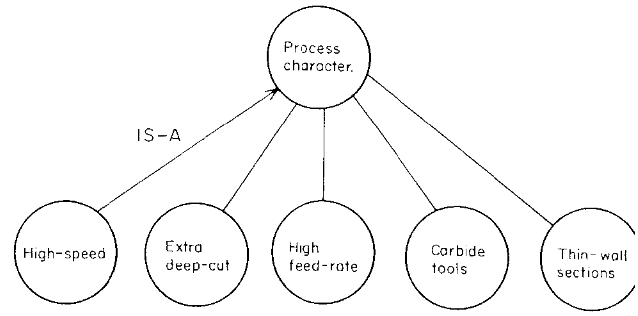

FIGURE 11. Process characteristics

{{**RUST**
IS-A: FIXED-PROCESS)DIAGNOSTIC
ACTIVATES: (NO-RUST 1.0)
IMPORTANCE: 10
STATUS:}}

A 'change' process diagnostic results in the alteration of one or more fundamental cutting fluid properties (Figure 13).

{{**LUBRICITY-DIAG**
IS-A: CHANGE-PROCESS-DIAGNOSTIC
ACTIVATES: (TOOL-LIFE LUBRICITY 1.0)
STATUS:}}

An example of a 'shop requirement' is the restriction of the cutting fluid selections to only soluble fluids:

{{**SOLUBLES-ONLY**
IS-A: SHOP-REQUIREMENT
ACTIVATES: (SOLUBLES 1.0)
STATUS:}}

4.3 PROPERTY REPRESENTATION

The diagnosis phase of GREASE generates a set of constraints. These constraints restrict the values of the properties of the cutting fluid to be chosen. Some cutting fluid properties, such as ODOR-MASKANT

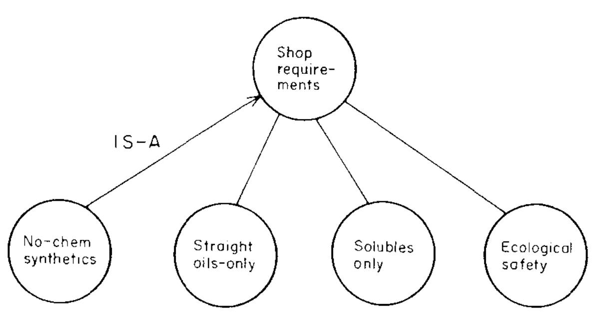

FIGURE 10. Shop requirements

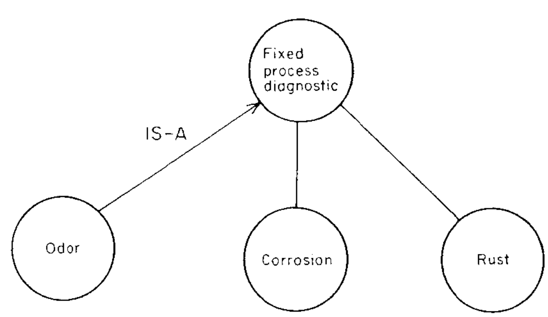

FIGURE 12. FIXED process diagnostics

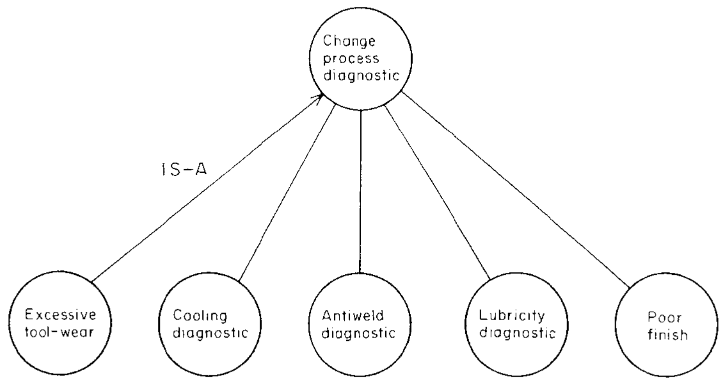

FIGURE 13. CHANGE process diagnostics

PROPERTY, are simple properties whose values can be directly measured and are independent of any of the other properties. However, the values of some properties, such as lubricity and antiweld, are functions of the values of other properties. Representing cutting fluid properties as schemata allows the functional relationships between these properties to be easily represented.

Cutting fluid properties in GREASE are represented both as slots in the **CUTTING-FLUID** schema (Figure 9) and by separate schema, as instances of the **PROPERTY** schema (Figure 14).

{{**PROPERTY**
PRE-REQUISITE: T
comment: Must evaluate to true in order for property to have a non-nil value. This can be used to restrict the applicability of properties.
UNITS:
VARIES:
comment: The properties that are a function of the given property.
VARIES-WITH:
comment: The properties that the given property is a function of.
CHANGED-BY:
comment: A list of goals for this property.
IS-A + INV: FINISH TOOL-LIFE CUTTING-FLUID-PROP
IMPORTANCE:
comment: How important the property is to the customer. This applies only to top-level properties such as TOOL-LIFE and FINISH.
VALUE-FOUND-IN:
comment: The name of the slot in the cutting fluid schema where the value of this property is stored, if it directly measured; otherwise it is nil.
CHANGE-MEASURE: MEASURE-CHANGE
comment: The function which calculates the value of the given property in terms of the values of lower level properties.}}

FIGURE 14. The property schema

The cutting fluid properties are represented both as a taxonomy depicting classes of properties (Figure 15) and as a dependency network based upon the functional dependence of the properties upon each other. A cutting fluid property which is functionally dependent upon another cutting fluid property is said to 'vary-with' that property. This 'varies-with' hierarchy is depicted in Figure 16.[7]

As an example, an instance of the property **TOOL-LIFE** is examined more closely.[8] **TOOL-LIFE** is an estimate of how well a given cutting fluid will reduce the cost of tool replacements in a given shop; this value is represented as the value of a slot in each cutting fluid schema. Its functional dependence upon other properties is represented by the 'varies-with' relations of the **TOOL-LIFE** schema to other schemata.

{{**TOOL-LIFE**
IS-A: PROPERTY
VARIES-WITH: LUBRICITY COOLING ANTIWELD VISC
CHANGED-BY: NIL
comment. Since there are no 'change' goals which directly affect tool life, the value is NIL.
DERIVES: T44 T43
comment: A list of 'change' goals which indirectly affect tool life by changing the properties tool life 'varies-with'.
VALUE-FOUND-IN: NIL
comment: Since this is always a calculated property, rather than a measured property, the value is nil.
VALUE: NIL}}

[7] The deepest level of properties of the 'varies-with' hierarchy is not specified here for proprietary reasons.

[8] It should be noted that the value of the tool-life property is actually dependent upon which process the cutting fluid is applied to. This process dependency is indicated in the individual change-goal schema listed in the DERIVES slot.

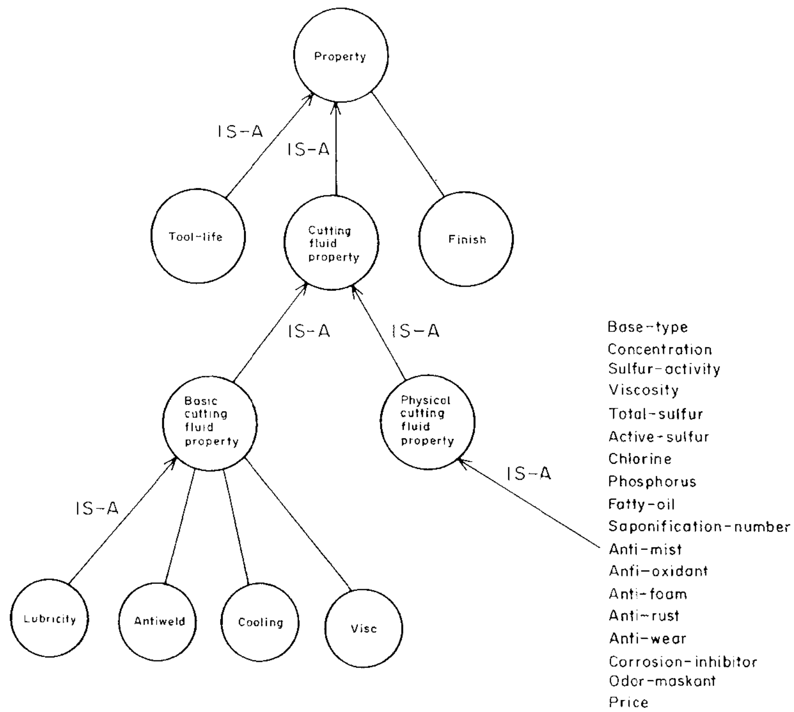

FIGURE 15. Cutting fluid property taxonomy

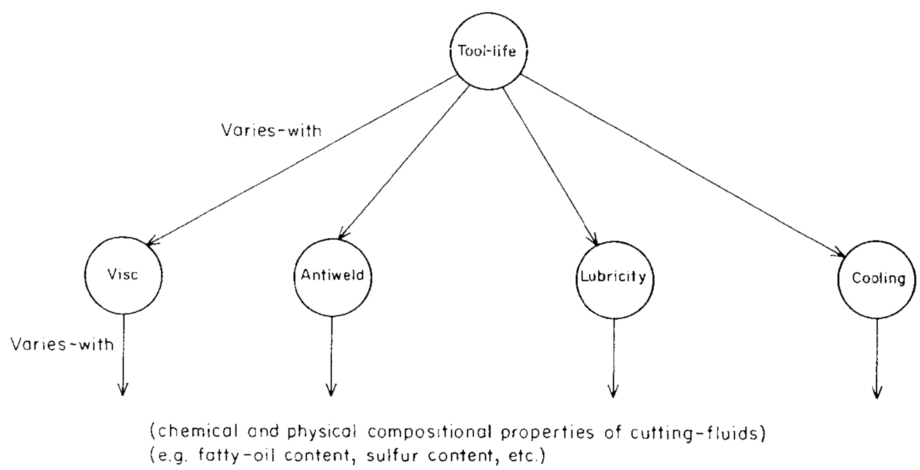

FIGURE 16. The 'varies-with' hierarchy of cutting fluids

4.4 GOAL NETWORK

Constraints within GREASE affect the selection and rating of the cutting fluids through 'goals' to the fluid evaluator. The 'goals' are posted to satisfy 'symptoms' e.g. diagnostics, atypical operating conditions, and user requirements, identified during the 'shop definition' phase of GREASE. 'Goals' also satisfy 'operation' and 'material' constraints such as the 'no-active-sulfur' constraint for cast aluminum.

4.4.1 Goal types

There exist two types of 'goals' within GREASE:

FIXED GOALS—These goals reflect condition which 'must' be satisfied for all processes during th fluid evaluation. If any "fixed goals" that ar posted fail for a particular fluid being evaluatec that fluid is rejected.

CHANGE GOALS—these goals require a chang in a cutting fluid property from some current o 'ideal' starting value. 'Change goals' affect th properties of lubricity, cooling, antiweld, and viscosity and are posted for the individual machining processes affected in the shop.

The **GOAL** schema is the root of the goal network:

```
{{GOAL
  IN-A + INV: CHANGE-GOAL FIXED-GOAL
  IMPORTANCE:
  STATUS:}}
```

The STATUS slot is set to a *process-id* list for those processes posting a particular CHANGE goal; otherwise, *nil.* The IMPORTANCE slot, not currently used, is reserved to allow a prioritization of goals.

'Fixed goals' are instances of the **FIXED-GOAL** schema. An example of a fixed goal is **NO-CHLORINE**:

```
{{NO-CHLORINE
  IS-A: FIXED-GOAL
  TEST: NO-CHLORINE-TEST
  IMPORTANCE: 10
  STATUS:}}
```

The TEST slot contains a predicate function which determines if the goal is met for the specific fluid being tested.

'Change goals' are instances of the **CHANGE-GOAL** schema:

```
{{CHANGE-GOAL
  IS-A: GOAL
  CHANGES:
  DERIVED-FROM:
  GOAL-OF:
  DEGREE: 0.0
  GOAL-FUNC: MAINTAIN-HIGH}}
```

Each "change goal" is specific to a particular machining process and is generated dynamically as it is being posted. The slots are defined as follows:

CHANGES—a relation indicating the fundamental cutting fluid property affected by the goal (i.e. cooling, lubricity, antiweld or viscosity);

DERIVED-FROM—a relation indicating the higher-level property (such as tool-life or finish) that the goal attempts to optimize;

GOAL-OF—the machining process that the goal is affecting;

DEGREE—the absolute value of the fundamental property resulting from the goal (i.e. the target value of the property for the process);

GOAL-FUNC—contains the name of a function which compares the value of the fundamental property for this goal against a tested fluid and returns a fluid rating.

4.4.2 *Goal activation*

Goals are activated by processing posted constraints originating from materials, operations, and symptoms identified in the shop requirements phase. The ACTIVATES slot is used to identify the constraints.

An example of a FIXED goal is the 'no-sulfur' constraint of cast-aluminum:

```
{{ALUMINUM-CAST
  IN-GROUP: GROUP7A
  USE-OF: ALUMINUM
  ACTIVATES: (NO-SULFUR 10)}}
```

FIXED goals are specified as two element lists—the first element is the name of the schema representing the FIXED goal, and the second element optionally contains a value used by the test function to indicate, for example, the maximum acceptable percentage level of the property.[9] When multiple goals exist for a constraint, they are appended to the ACTIVATES slot.

CHANGE goals are specified by three element lists as in the example of high-speed:

```
{{HIGH-SPEED
  IS-A: PROCESS-CHARACTERISTIC
  ACTIVATES: (TOOL-LIFE COOLING 1.0)
    (TOOL-LIFE ANTIWELD 1.0)
  STATUS:}}
```

The first element specifies the 'high-level' property being optimized by the goal (the current GREASE implementation always specifies TOOL-LIFE). The second element specifies the cutting fluid property whose alteration will satisfy the posted constraint. The third element represents the amount that the cutting fluid property will be altered (indicated as a multiplier to a 'typical change' of the property, as discussed in the next section).

4.5 EMPIRICAL SELECTION CAUSAL NETWORK

Part of the evaluation function generation stage is the identification of an ideal fluid for each material-process pair. The ideal fluid provides a starting point for the chemical properties of the optimal fluid. (The properties of the ideal are altered according to the known constraints.) The ideal fluids are organized as a matrix, indexed by material group and operation. The network embodies the results of using cutting fluids in actual machine operations and,

[9] The *no-chlorine* goal has a second argument of '1.0', indicating that a fluid containing up to 1% chlorine can still pass the goal. FIXED goals that do not use the second argument have been arbitrarily set to a value of 10.

therefore, represents a wealth of experimental knowledge. The network is augmented by expert system knowledge of how to satisfy fundamental property requirements in a selection situation in terms of chemistry alterations which affect these properties.

4.5.1 *Material/operation ideal fluid table*

The network is represented in terms of individual schemata assigned to specific 'processes' in GREASE. There exist a schema for each possible combination of operation severity and material machinability group. Each schema accomplishes the following:

- represents the 'Ideal Fluid' in terms of cutting fluid properties for the specified maching process;
- contains property values representing the granularity of change for a "change goal";
- specifies rating functions for each property when evaluating actual fluids;
- specifies sensitivities of each property for the specified process.

Each schema in the network is an instance of **TBL** schema:

```
{{TBL
  SATISFIES-SEVERITY-OF:
  SATISFIES-MACHINABILITY-OF:
  INSTANCE + INV: TBL11 TBL12 TBL13 . . .
  OIL-NAME:
  COOLING: 50.0
  LUBRICITY: 40.0
  ANTIWELD: (190.0 STANDARD-CHANGE)
  ACTIVITY: 230.0
  VISC: 3.5
  TOOL-LIFE:
  FINISH:
  RECOMMENDATIONS:}}
```

The slot SATISFIES-SEVERITY-OF denotes the operation to which severity referred; SATISFIES-MACHINABILITY-OF designates the machinability class of the material. The slot OIL-NAME specifies the name of an 'ideal' cutting fluid with fundamental properties that best satisfy the process.

The slots COOLING, LUBRICITY, ANTIWELD, ACTIVITY and VISC contain values that represent the magnitude of a 'typical change' or increment in the properties if atypical operating conditions are used or a machining diagnostic is observed. These same values also specify the width of the utility functions used by the fluid evaluator.[10] If the slot values are represented as a list, the last list element represents the utility function to be used to rate the specified property by the fluid evaluator.[11]

The slot TOOL-LIFE contains a meta-slot attachment that contain facets representing the 'sensitivity' of each fundamental cutting fluid property to affecting 'tool-life' for the process. The slot FINISH is intended to be used similarly in a future expansion of GREASE.

An example of a schema associated with a process of operation severity 'severity1' and material machinability 'group2' is below:

```
{{TBL12
  INSTANCE: TBL
  SATISFIES-SEVERITY-OF: SEVERITY1
  SATISFIES-MACHINABILITY-OF: CLASS2
  OIL-NAME: FERROUS12
  COOLING: 25.0
  LUBRICITY: 100.0
  ANTIWELD: (150.0 STANDARD-CHANGE)
  VISC: 3.5
  TOOL-LIFE:
    INSTANCE: TL1
  FINISH:
    INSTANCE: FINISH1}}
```

'Ideal' cutting fluids in GREASE are hypothetical fluids with property values corresponding to optimal selections for the machining processes embodied in the empirical selection causal network. The property values in the 'ideal' fluids have been tuned by cutting fluid experts. These fluids are used by GREASE to determine a starting chemistry for a specified process before posting any goals.

The 'ideal' fluid is a function of the operation severity and machinability determined by the material-class. In addition, a distinction is made between 'ideal' fluids for ferrous and non-ferrous materials. As a result, there are forty 'ideal' fluids for operation severities 'severity1' through 'severity10' and ferrous material 'class1' through 'class4'. The naming convention is 'FERROUSmn' where 'm' represents the operation severity and 'n' represents the material class.

```
{{FERROUS11
  INSTANCE: IDEAL-FLUID
  LUBRICITY: 100.0
  ANTIWELD: 300.0
  COOLING: 600.0
  KINEMATIC-VISCOSITY: 38.6}}
```

Similarly, there are fifty 'ideal' fluids for the non-ferrous materials corresponding to material

[10] If the slot contains two values, the first represents the utility function width parameter, and the second represents the diagnostic increment parameter.

[11] If the utility function is unspecified, the 'default' utility function, MAINTAIN-HIGH, will be used to rate the property by the fluid evaluator.

classes 'class0' through 'class4' and operation severities 'severity1' through 'severity10'. (Note: The empirical selection causal network contains 100 machining processes since there are ten material groups and ten severity classes. There are only 90 'ideal' fluids since 'group5a' and 'group5b' share the same material class 'class0'.)

4.5.2 Sensitivities

The sensitivities, as previously mentioned, refers to the ability that each fundamental cutting fluid property, such as lubricity, has to influence tool life or finish. The sensitivities are represented as meta-slots associated with either the TOOL-LIFE or FINISH slot of an instance of the **TBL** schema. The meta-slots themselves are instances of the **M-TOOL-LIFE** schema:

```
{{M-TOOL-LIFE
  IS-A + INV: TL10 TL9 . . . TL3 TL2 TL1
  LUBRICITY: 2.0
  ANTIWELD: 2.0
  COOLING: 2.0
  VISC: 0.0
  ACTIVITY: 0.0}}
```

The **TL1** schema is the meta-slot of the **TBL12** TOOL-LIFE slot in the previous example:

```
{{TL1
  IS-A: M-TOOL-LIFE
  LUBRICITY: 0.77
  ANTIWELD: 4.62
  COOLING: 0.77
  VISC: 3.85
  ACTIVITY: 0.0}}
```

The slot values, or facets represent the individual property sensitivities. The values were determined from interviewing the experts as to how significant, or important, each of the properties was for each machining operation. Each property was rated according to a six-valued ordinal scale ranging from 'low' to 'very-high'. The ratings were then converted to the numeric scale 1–6, and the sum of the ratings for a fluid was normalized to 10.0.[12]

4.5.3 Utility functions

The utility functions are rating functions used by the fluid-elevator to determine how well a 'candidate' fluid from the product line or blend matches the designed 'optimal' fluid. The shape of these functions is based both upon cutting fluid property behavior and empirical testing. The significant functions depicted in Figure 17 are **maintain**, **maintain-high**, **maintain-low**, and **standard-change**. The properties of these functions are as follows:

maintain—used when a property must be restricted to a narrow range, such as 'maintaining' a viscosity value in 'deep-drilling'.

maintain-high—used as the default, returns a low rating when the fluid has a value less than the 'optimal', but a maximum rating for higher values. Cooling and lubricity generally have this property.

maintain-low—returns a high rating when the desired property value is below or equal to the optimal, but a low rating when it exceeds the optimal.

standard-change—generally used by antiweld, returns a low rating if the value of the fluid is both lower than the optimal, and greater than a 'typical change' higher than the optimal value. The tool-life drops off as antiweld is increased beyond a reasonable value for a machining process.

The functions are variations of a mathematical 'normal' curve. The curve width is determined by the 'typical change' value for the property in the empirical selection causal network and corresponds to the variance in the 'normal' curve formula.[13] The functions return a rating with values between 0.0–1.0.

The 'typical changes' have been determined by interviewing the cutting fluid experts regarding how much a property must *typically* change to satisfy an upset resulting in a diagnostic symptom or to satisfy an atypical operating condition—this can be interpreted as the granularity of the property change. Examples of typical changes include the amount that 'cooling' must be changed to satisfy a 'blued chips' diagnostic, or to correct for a 'high speed' operation. Multiples of these 'typical changes' are used by GREASE to satisfy posted goals.

5. Implementation of GREASE

Consider the following situation:

> Select a cutting fluid for the machining of three materials: titanium, which is 30% of the material machined; B-1111, which is 50% and GROUP3, which is 20%. The tapping process is performed on GROUP3 material 100% of the

[12] The ACTIVITY slot is not utilized in the current GREASE prototype. Also, the meta-slots **TL1**→**TL10** have been initialized to the same sensitivities for the different machinability groups of each operation severity. This is only an approximation. Tuning of GREASE will reveal distinct sensitivities which are a function of both operation severity and material machinability. In this case, specific meta-slots, such as **TL12**, attached to slot TOOL-LIFE of schema **TBL12**, may be added to GREASE to reflect this.

[13] In tuning GREASE, it was found that a curve width approximately 1.5 × 'typical change' for a property symptom results in GREASE selections in better agreement with the experts in most cases.

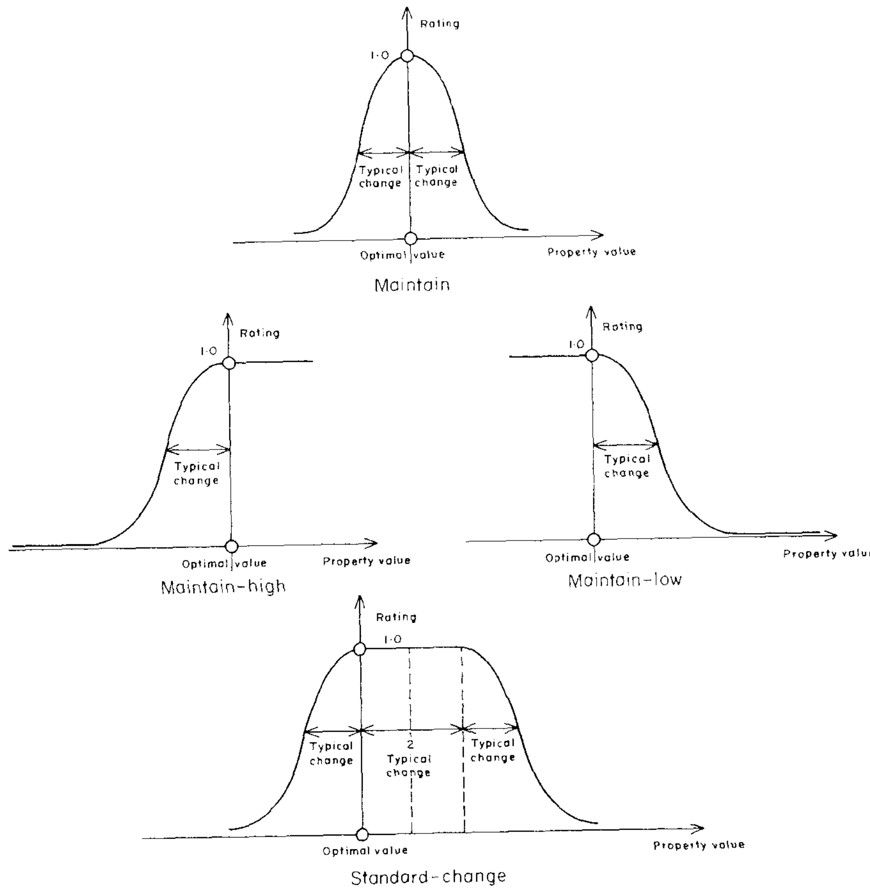

FIGURE 17. The four utility functions

time, automatic multiple spindle machining is performed on titanium 30% of the time and on B-1111 50% of the time. The problems which have arisen for the tapping of GROUP3 are: discolored tool edges, large built up edge and poor finish. Automatic multiple spindle of titanium results in excessive tool wear, and on B-1111 it is a high speed operation.

In this section, each processing stage of GREASE is described as it selects the best fluid for this situation.

5.1. SHOP DEFINITION

The first stage extracts from the user a definition of the problem, which includes information about the machining operations performed, the materials to be machined, the processes relating operations and materials, the current cutting fluid used, and additional customer requirements. The information is extracted through the asking of simple questions.

; OPERATION SPECIFICATION
Machine operation (if not in system, closest type)?
  automatics-multiple-spindle
Brass, bronze or copper in machinery (y/n)? n
Machine operation (if not in system, closest type)? tapping
Brass, bronze or copper in machinery (y/n)? n

; MATERIAL SPECIFICATION
AISI #, group #, or material? B-1111
% of all materials? 50
AISI #, group #, or material? titanium
% of all materials? 30
AISI #, group #, or material? group3
% of all materials? 30
GROUP3 can't exceed: 20.0%
% of all materials? 20

; PROCESS SPECIFICATION
Material ? ?
Possible responses:
B-1111
GROUP3
TITANIUM
Material ? 1111
Operation ? auto
% of this material ? 100
Importance of maintaining tool-life [1–10] 1?
High speed (y/n) ? y
Extra deep cut (y/n) ?

High feed rate (y/n) ?
Carbide tools (y/n) ?
Thin wall sections (y/n) ?
Excessive tool wear (y/n) ?
Long chips (y/n) ?
Discolored tool edges (y/n) ?
Hot work pieces (y/n) ?
Smoke (y/n) ?
Chip welding (y/n) ?
Large built-up edge (y/n) ?
Tool seizure (y/n) ?
Poor or sub-standard finish (y/n) ?
Rust (y/n) ?
Corrosion (y/n) ?

; the same questions are repeated for GROUP 3 and Titanium.

; CURRENT FLUID SPECIFICATION
Current cutting fluid name ? 31B

; ADDITIONAL REQUIREMENTS
ecological safetype required (y/n)?
Recommend only straight oils (y/n)?
Recommended only soluble oils (y/n)?
Can synthetics be recommended (y/n)?

; SHOP SUMMARY
OPERATIONS
TAPPING
AUTOMATICS-MULTIPLE-SPINDLE
MATERIALS
GROUP3 20%
TITANIUM 30%
B-1111 50%
PROCESSES
(TAPPING GROUP3)
material%: 100
process%: 20.0
tool-life-importance: 1.0
DISCOLORED-TOOL-EDGES
LARGE-BUILT-UP-EDGE
POOR-FINISH

(AUTOMATICS-MULTIPLE-SPINDLE TITANIUM)
material %: 100
process %: 30.0
tool-life-importance: 1.0
EXCESSIVE-TOOL-WEAR

(AUTOMATICS-MULTIPLE-SPINDLE B-1111)
material %: 100
process %: 50.0
tool-life-importance: 1.0
HIGH-SPEED
CURRENT FLUID
31B
Interpretation of Diagnostics
In Process: (TAPPING GROUP3)
  Discolored tool edges
    ⇒ Cooling Problem.
  Large built-up edge
  Poor Finish
    ⇒ Antiweld Problem.
In Process: (AUTOMATICS-MULTIPLE-SPINDLE TITANIUM)
  Excessive tool wear
    ⇒ Cooling Problem.
  excessive tool wear
    ⇒ Lubricity Problem.
  Excessive tool wear
    ⇒ Antiweld Problem.

Based on this dialogue, GREASE creates a model of the shop. The automatic multiple spindle operation is represented as:

{{**T33**
INSTANCE: OPERATION-SPEC
OPERATION-IS-A: AUTOMATICS-MULTIPLE-SPINDLE
STATUS: T
WORN-MACHINERY: NIL}}

This schema is characterized by two slots OPERATION-IS-A and WORN-MACHINERY. Each instance points to an element of the operation's taxonomy by means of the inherited slot OPERATION-IS-A. The slot WORN-MACHINERY is currently used to indicate if the machine contains brass or bronze components which would be affected by active sulfur in the cutting fluid.[14]

The material specification of titanium representing 30% of the material machined is represented as:

{{**T34**
INSTANCE: ANONYMOUS-MATERIAL-SPEC
MATERIAL-IS-A: TITANIUM
MATERIAL-PERCENTAGE: 30}}

Materials can be specified by their AISI number or machinability group. A material specified by its AISI number will be an instance of the **AISI-MATERIAL-SPEC** schema; otherwise, it will be an instance of the **ANONYMOUS-MATERIAL-SPEC** schema. Each instance will point to an element of the materials taxonomy by means of the inherited slot MATERIAL-IS-A. The slot MATERIAL-PERCENTAGE, representing the percentage of the specified material in the shop, is used as a measure of the importance of the specified material among all the materials machined in the shop.

The specification that all titanium is machined on a multi-spindle machine with the diagnostic of excessive

[14] Originally WORN-MACHINERY indicated that the machine was worn and contained brass and bronze components. Cutting fluid would seep into these machines causing corrosion.

tool wear is defined by a **PROCESS-SPEC** schema:

```
{{T36
  INSTANCE: PROCESS-SPEC
  INVOLVES-MATERIAL: T34
  INVOLVES-OPERATION: T33
  OPERATION-PERCENTAGE: 100
  IMPORTANCE: 1
  PROCESS-PERCENTAGE: 30.0
  TABLE-POSITION: TBL57B
  NAME: (AUTOMATICS-MULTIPLE-SPINDLE
    TITANIUM)
  HIGH-SPEED: NIL
  EXTRA-DEEP-CUT: NIL
  HIGH-FEED-RATE: NIL
  CARBIDE-TOOLS: NIL
  THIN-WALL-SECTIONS: NIL
  EXCESSIVE-TOOL-WEAR: T
  LONG-CHIPS: NIL
  DISCOLORED-TOOL-EDGES: NIL
  HOT-WORK-PIECES: NIL
  SMOKE: NIL
  COOLING-DIAG: NIL
  SOFT-DRAGGY-METAL: NIL
  LUBRICITY-DIAG: NIL
  CHIP-WELDING: NIL
  LARGE-BUILT-UP-EDGE: NIL
  TOOL-SEIZURE: NIL
  ANTIWELD-DIAG: NIL
  POOR-FINISH: NIL
  RUST: NIL
  CORROSION: NIL}}
```

This schema is characterized by the three slots INVOLVES-MATERIAL, INVOLVES-OPERATION and OPERATION-PERCENTAGE. Each instance points to an instance of the OPERATION-SPEC schema and an instance of the MATERIAL-SPEC schema, by means of the inherited slots INVOLVES-OPERATION and INVOLVES MATERIAL, respectively. The slot OPERATION PERCENTAGE records the proportion of the specified operation among all the operations performed on the corresponding material. The slo IMPORTANCE holds a number representing the importance of maintaining the tool-life corresponding to the process. A default value has been set to 1. The slo PROCESS-PERCENTAGE records the percentage of the specified process relative to all machining processes in the shop. This process-percentage is internally computed as follows

$$Process\ percentage = op \times mat$$

where *op* represents the value of the slo OPERATION-PERCENTAGE, and *mat* is the value of the slo MATERIAL-PERCENTAGE of the corresponding instance o the **MATERIAL-SPEC** schema. TABLE-POSITION hold a pointer to a table entry of the empirical selectio causal network that represents the process. NAMI holds a string corresponding to the name of the process.

The slots HIGH-SPEED, EXTRA-DEEP-CUT, HIGH-FEED RATE, CARBIDE-TOOLS and THIN-WALL-SECTIONS contai the operation characteristics of the process.

In a similar fashion, the slots EXCESSIVE-TOOL-WEAR LONG-CHIPS, DISCOLORED-TOOL-EDGES, HOT-WORK-PIECES SMOKE, SOFT-DRAGGY-METAL, CHIP-WELDING, LARGE BUILT-UP-EDGE, TOOL-SEIZURE, POOR-FINISH, RUS

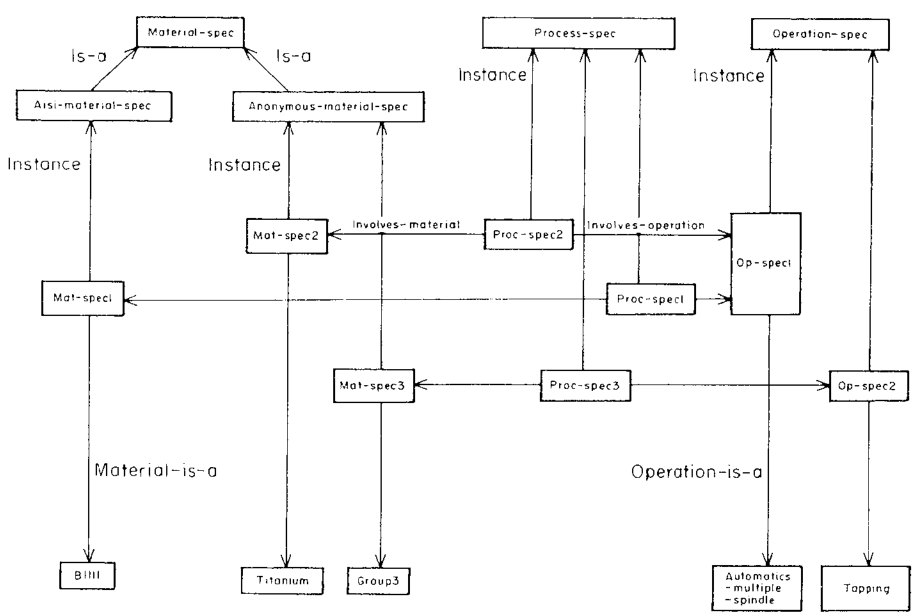

FIGURE 18. Shop network for example

and CORROSION record the various diagnostics related to the process. The slots COOLING-DIAG, LUBRICITY-DIAG and ANTIWELD-DIAG contain the results of diagnostic analysis of symptoms that indicate imbalances in cutting fluid cooling, lubricity, and antiweld.

Figure 18 is an example that illustrates how the shop definition specifications just introduced combine with each other to create a network description of the shop:

proc-spec1 represents a multiple-spindle automatic operation on **b1111** steel and has a value of 'T' for the HIGH-SPEED slot.

The current fluid specification identifies the currently used cutting fluid in the shop. The current fluid specification is necessary if diagnostic information is to be used for the shop. The current fluid chemistry allows GREASE to determine what cutting fluid property levels resulted in the specified diagnostics, and what final property levels will be necessary to treat the diagnostics. The current fluid is also used to assist GREASE in determining levels of cutting fluid properties to maximize tool-life at minimum cost.

The specified fluid '31B' is in the product line. GREASE will represent this specification in the **CURRENT-FLUID** schema as an instance of '31B':

```
{{CURRENT-FLUID
  INSTANCE: 31B
  NAME: 31B
  TYPE: OIL
  DILUTION-RATIO:
  EMULSIFIER-PERCENTAGE:
  KINEMATIC-VISCOSITY:
  TOTAL-SULFUR-PERCENTAGE:
  ACTIVE-SULFUR-PERCENTAGE:
  CHLORINE-PERCENTAGE:
  FATTY-OIL-PERCENTAGE:
  COOLING:
  LUBRICITY:
  ANTIWELD:
  ACTIVITY:}}
```

The NAME slot contains the name of the user-specified current fluid. If the current fluid is in the product line, the schema CURRENT-FLUID is specified to be an instance of that corresponding product in the cutting fluid taxonomy. The chemical properties of the current fluid are then inherited via the INSTANCE link. If the name refers to a soluble oil, a dilution ratio is prompted for.

If the current fluid is not in the product line, the user is prompted for the entire chemical composition. The cooling, lubricity, antiweld, and activity levels of the cutting fluid are then internally computed according to the values of the chemical composition slots.[15] These properties are then stored in the COOLING, LUBRICITY, ANTIWELD and ACTIVITY slots.[16] The chemical composition of the fluid is characterized by the slots: KINEMATIC-VISCOSITY, TOTAL-SULFUR-PERCENTAGE, ACTIVE-SULFUR-PERCENTAGE, CHLORINE-PERCENTAGE, FATTY-OIL-PERCENTAGE. The slot TYPE holds the type of the fluid *oil* or *soluble*. In case of a water-based fluid, the slot DILUTION-RATIO will be set to the corresponding value.

Diagnostics encountered with the current-fluid are recorded by the various symptom slots corresponding to the instances of the **PROCESS-SPEC** schema.

5.2 DIAGNOSIS AND EVALUATION FUNCTION GENERATION

```
Activation of Goals
OPERATIONS CHARACTERISTICS
  No goals posted . . .
MATERIALS CHARACTERISTICS
TITANIUM
  satisfying goal > NO-CHLORINE
USER PREFERENCES
No goals posted . . .
FIXED PROCESS DIAGNOSTICS
  No goals posted . . .
PROCESS CHARACTERISTICS
HIGH-SPEED
  In (AUTOMATICS-MULTIPLE-SPINDLE B-1111)
    increases COOLING by 50.0 to 645.0
  In (AUTOMATICS-MULTIPLE-SPINDLE B-1111)
    increases ANTIWELD by 100.0 to 170.0
CHANGE PROCESS DIAGNOSTICS
POOR-FINISH
  In (TAPPING GROUP3) maximizes 31B ANTIWELD
    by 177.0 to 404.0
ANTIWELD-DIAG
  In (TAPPING GROUP3) maximizes 31B ANTIWELD
    by 177.0 to 404.0
COOLING-DIAG
  In (TAPPING GROUP3) maximizes 31B COOLING BY
    50.0 to 778.429993
EXCESSIVE-TOOL-WEAR
  In (AUTOMATICS-MULTIPLE-SPINDLE
    TITANIUM) maximizes 31B LUBRICITY by 30.0 to
    64.0
  In (AUTOMATICS-MULTIPLE-SPINDLE
    TITANIUM) maximizes 31B COOLING by 50.0 to
    838.799988
```

[15] The empirical formulas correlating cutting fluid properties with chemical composition are proprietary and not described here.

[16] GREASE cannot calculate the cutting fluid properties of 'chemical synthetic fluids' since the effect of different chemical species in these fluids relative to the cutting fluid properties is unknown. GREASE will reject these fluids as invalid current fluids.

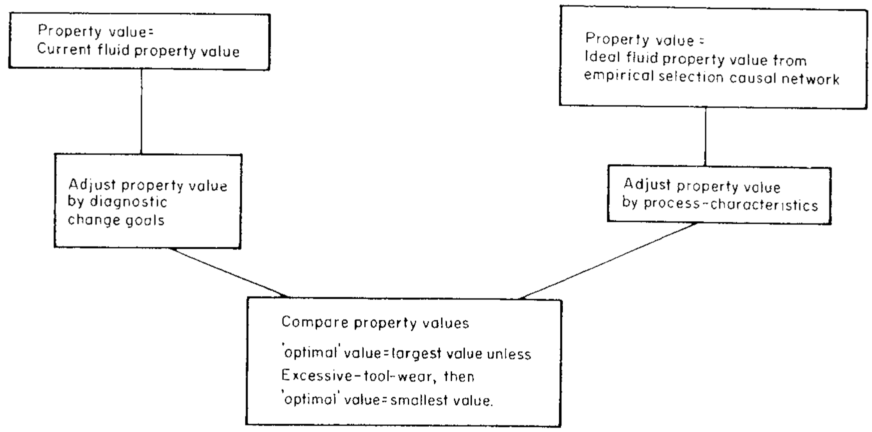

FIGURE 19. 'Change' goal resolution

In (AUTOMATICS-MULTIPLE-SPINDLE TITANIUM) maximizes 31B ANTIWELD by −100.0 to 167.0

The next stage of recommendation generates cutting fluid goals from diagnostic symptoms, atypical operating conditions, and operation and material constraints. Constraints are used by GREASE to convert the 'ideal' fluids for each shop process to 'optimal' fluid property goals. Goals are generated to satisfy the requirements of a particular machining process such as treating diagnostics, satisfying operator preferences, and correcting for exceptional machining conditions.

Fixed goals are generated by examining operation and material constraints user preferences and process diagnostic symptoms. Change goals are generated by identifying symptoms such as process characteristics (i.e. 'high-speed', 'deep-cut', etc.) and process diagnostics (i.e. 'built-up-edge', 'poor-finish', etc.) and creating a goal for each fundamental cutting fluid property (i.e. cooling, lubricity, antiweld, viscosity) for each process identified in the 'shop-definition' phase. These goals are primed with fundamental property values corresponding to the 'ideal fluid' for the machine process.

Constraints resulting from process characteristics such as 'high-speed' and 'deep-cut' are 'additive' requiring a goal associated with a machine process to sum the contributions from all identified process characteristic symptoms for the affected property. The sum is then added to a 'reference point' for the property to create an absolute 'optimal' value (i.e., a goal), which, if met, will satisfy the symptoms. The 'reference point' is derived from either a 'current fluid' identified in the shop definition phase, or the 'ideal fluid' indicated in the 'empirical selection causal network'.

Constraints from process diagnostics are 'maximizing', where the affected property value is the maximum value for any identified diagnostic. Process diagnostic generated goals are only posted if a 'current fluid' is indicated that serves as a reference for which the goals can improve upon.

Goal generation attempts to optimize the cutting fluid recommendation wherever possible (Figure 19). It does this by using the property values of the 'ideal fluid' in its reference point selection, rather than the 'current fluid' values which may be an *improperly recommended fluid.* If diagnostics are posted, GREASE uses the 'current-fluid' as a reference point upon which to maximize the diagnostic goals, but then compares the results with the corresponding 'ideal fluid' property, adjusted by process characteristic goals, and selects the maximum value to better correct the diagnostic conditions.[17]

In the example, titanium generated a 'fixed-goal' of *no-chlorine,* and the excessive-tool-wear diagnostic generated maximized 'change goals' for cooling, lubricity, and antiweld. Since a current-fluid was specified '31B', the cooling, lubricity, antiweld, and visc property levels for the initial optimal fluid for the process were primed with the corresponding '31B' fluid values.

The 'change-goals' that were generated for this example are illustrated as in Figure 20. Schemata 'goal1, goal2, goal3, goal4' specify, in the DEGREE slot, the optimal property values for lubricity, cooling, and visc for the designated process.

[17] If 'excessive tool wear' was indicated, the minimum value is selected.

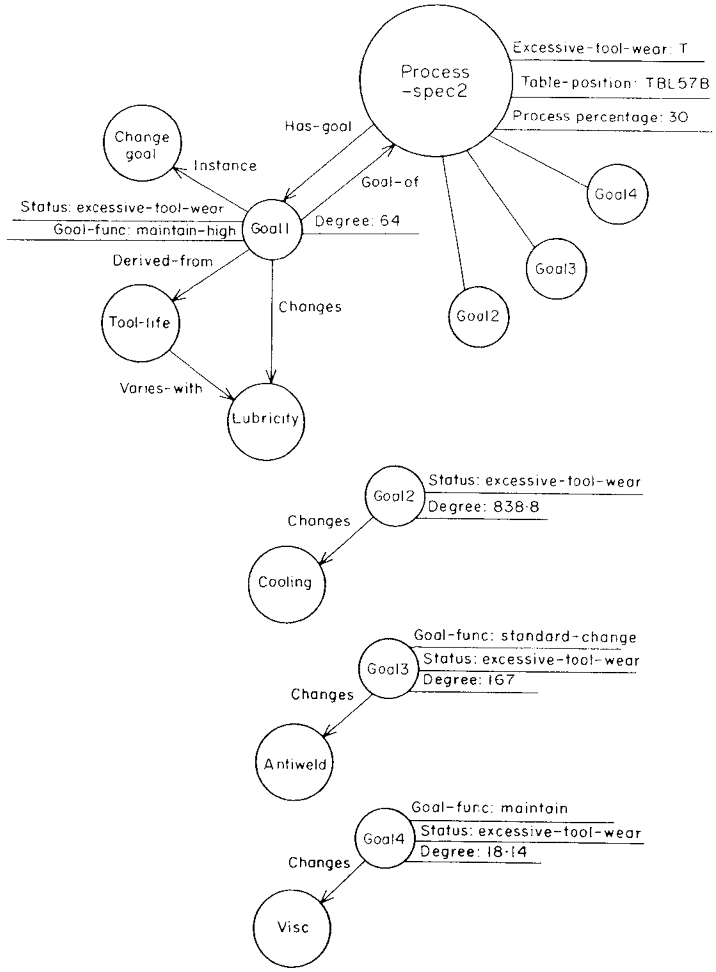

FIGURE 20. Goal generation for (AUTOMATICS-MULTIPLE-SPINDLE TITANIUM) process

5.3 CANDIDATE FLUIDS GENERATION

The candidate fluids correspond to the search space used by GREASE from which a recommendation is determined.

The candidate fluids are selected from the cutting fluid product line. Six dilution ratios for each soluble fluid have been determined along with the property values corresponding to these dilutions. Each dilution ratio determines a separate product. If a current fluid has been specified, it is included to allow it to be rated relative to the product line.

5.4 FLUID EVALUATION

```
TESTED-FLUID⇒31A
  TOOL-LIFE: 8.729034 (TAPPING GROUP3)
  TOOL-LIFE: 8.893048 (AUTOMATICS-MULTIPLE-
    SPINDLE TITANIUM)
  TOOL-LIFE: 9.866111 (AUTOMATICS-MULTIPLE-
    SPINDLE B-1111)
TESTED-FLUID⇒31B
  TOOL-LIFE: 7.877955 (TAPPING GROUP3)
  TOOL-LIFE: 8.800831 (AUTOMATICS-MULTIPLE-
    SPINDLE TITANIUM)
  TOOL-LIFE: 9.763288 (AUTOMATICS-MULTIPLE-
    SPINDLE B-1111)
TESTED-FLUID⇒41M
  TOOL-LIFE: 7.035658 (TAPPING GROUP3)
  TOOL-LIFE: 9.122429 (AUTOMATICS-MULTIPLE-
    SPINDLE TITANIUM)
  TOOL-LIFE: 9.602408 (AUTOMATICS-MULTIPLE-
    SPINDLE B-1111)
TESTED-FLUID⇒41D
  TOOL-LIFE: 6.142008 (TAPPING GROUP3)
  TOOL-LIFE: 7.537755 (AUTOMATICS-MULTIPLE-
    SPINDLE TITANIUM)
  TOOL-LIFE: 9.417368 (AUTOMATICS-MULTIPLE-
    SPINDLE B-1111)
TESTED-FLUID⇒41E
  NO-CHLORINE ⟨⟨failed⟩⟩
```

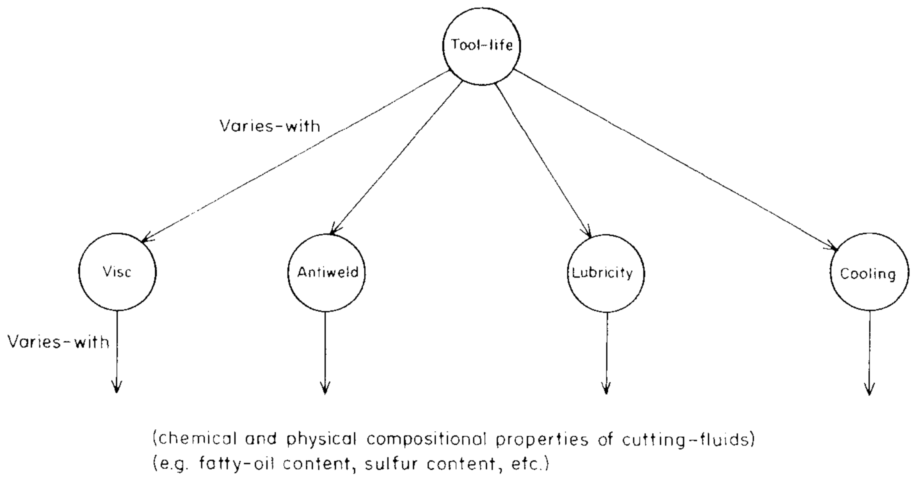

FIGURE 21. The 'varies-with' hierarchy of cutting fluids

Selection in GREASE is performed by a process known as fluid evaluation. Fluid evaluation attempts to optimize the value of the highest-level cutting fluid property for a machine shop—*tool-life*.[18] Optimal tool-life will result in minimal tool replacement costs, and optimized production.[19]

Fluid evaluation is accomplished by comparing the cutting fluid properties of each candidate fluid with the 'optimal' fluid for each specified machining process. An 'optimal' fluid is defined as one containing *ideal* property values to satisfy a specified machining process as determined by posting 'goals' against the 'ideal' fluid for a process. The relative approximation of a candidate fluid to each 'optimal fluid', as determined by utility functions, generates the overall rating of how well that fluid will perform.

The evaluation of a cutting fluid in terms of its properties is distinct for each cutting fluid property level (Figure 21) (Section 4.3).[20]

[18] The property *finish* is also a highest-level cutting fluid property. In the cutting fluid industry, however, optimal tool-life is of paramount importance, as long as finish is adequate. Consequently, optimal finish isn't attempted by the system. Inadequate finish is corrected for by a 'poor finish' diagnostic.

[19] Strictly speaking, or course, the customer wants to reduce overall costs, not only tool replacement costs, as much as possible. In order to estimate overall costs, however, it would be necessary to know the price of the cutting fluid. Since the price of a cutting fluid is a matter of negotiation between the salesman and the customer, this information is not known ahead of time. Therefore, we settle for giving information about the estimated costs of tool replacement, and leave it up to the salesman and customer to take into account the prices of the cutting fluids.

[20] The specific cutting fluid properties at the deepest level haven't been specified here for proprietary reasons.

5.4.1 *Evaluation of low level properties relative to a process*

Low level properties are properties whose values do not depend upon the values of any other properties. They are generally measured chemical properties such as the percentage of active sulfur, or physical properties, such as viscosity. The values of these properties are not calculated by GREASE. In the case of GULF products, these values are already known and stored in the GREASE knowledge base. In the case of non-GULF products, the values of these properties are prompted for (e.g. Current Fluid Specification).

5.4.2 *Evaluation of middle level properties relative to a process*

Middle level properties are properties whose values depend upon the values of other properties, and which have other properties whose values depend upon them. An example of such a property is lubricity. Its value depends upon the values of fatty content and viscosity; in turn, the values of tool life and finish depend upon the value of lubricity.

For efficiency, since the values of the middle-properties are functions of the chemistry of the cutting fluids, which very seldom change, they are pre-calculated and stored in the knowledge base. In the case of non-GULF products which haven't previously been encountered, the values of these properties are calculated from the values of the lower level properties during the current fluid specification.

5.4.3 *Evaluation of high-level properties relative to a process*

The single high-level property of cutting fluids to be evaluated is tool-life. It is measured on an arbitrary scale in which 1 is the worst value and 10 is the best value.

The tool-life of a cutting fluid is evaluated on the basis of how closely it matches a theoretical 'optimal' fluid for a process in terms of properties. The contribution of a particular property to the tool life value of a fluid depends upon two factors:

A *rating* measuring how closely the value of the property matches the 'optimal' value of that property.

A *coefficient* measuring how important that property is in determining the tool life known as the 'sensitivity' (Section 4.5.2).

The contribution of each property to the value of tool life is the product of the two factors listed above. The total tool life value is simply the sum of the contributions from each individual property.

To write the value of the tool life as an equation, let 'rating(property)' represent how closely the actual value of that property matches the 'optimal' value of the property. Then the equation for tool life is represented:

$$\begin{aligned}\text{Tool life} = a \times \text{rating(lubricity)} + b \times \text{rating(antiweld)} \\ + c \times \text{rating(cooling)} \\ + d \times \text{rating(viscosity)} \qquad (1)\end{aligned}$$

The determination of the rating and importance of a given property that tool-life depends upon will now be explained and illustrated with specific examples:

Importance—'sensitivity coefficient': Tool-life is evaluated with respect to a given process since the sensitivity coefficients are process dependent. For example, the tool life value of a given fluid depends upon how well the lubricity of the fluid matches the ideal lubricity value for that process; but how important it is for a fluid to match the ideal lubricity value is dependent upon what process the fluid is being applied to. In *easy to machine* metals, the lubricity value of a fluid is relatively unimportant; this is reflected in a low lubricity sensitivity coefficient (a low value of *a* in 1). In *difficult to machine* metals the lubricity value of a fluid is more important; this is reflected in a high value of the lubricity sensitivity coefficient (a high value of *a* in 1).

Since the sensitivity coefficients are process dependent, their values are stored in the empirical selection causal network which represents all possible processes that GREASE knows about. The empirical selection causal network schemata have a TOOL-LIFE slot. Attached to this slot is a meta-slot containing sensitivity coefficient facets for the middle-level properties upon which tool life depends, namely, COOLING, ANTIWELD, VISC and LUBRICITY. The sensitivity coefficient of viscosity in the equation for tool life is stored in the VISC slot, etc.

To illustrate these concepts, examine the process in Figure 20. The TABLE-POSITION slot indicates the schema TBL57B within the empirical selection causal network. Figure 22 illustrates the schema TBL57B. The tool-life meta-slot has a lubricity facet with a value of 0.77 representing the lubricity sensitivity coefficient for PROCESS-SPEC2.

Rating: In order to see how the rating of a property is calculated, consider how the optimal fluid is represented. A diagram of how an optimal fluid is represented was presented in Section 5.3. A goal schema is attached to each property that tool life "varies with". For example, **GOAL1** is the 'optimal' goal for **LUBRICITY** in process **PROCESS-SPEC2**. Each goal schema serves two major functions:

It contains the 'optimal' value that the attached property should have in order to maximize tool life for the process (DEGREE slot);

It specifies a function (GOAL-FUNC slot) which, when given an actual value of the property as input, returns a rating on a scale of 0 to 1 indicating how closely the actual value matches the desired value (Section 4.5.3).

For example, consider again **LUBRICITY** in Figure 20. The 'optimal' value of 'lubricity' to maximize tool life is stored in the DEGREE slot of **GOAL1** (i.e. 64).

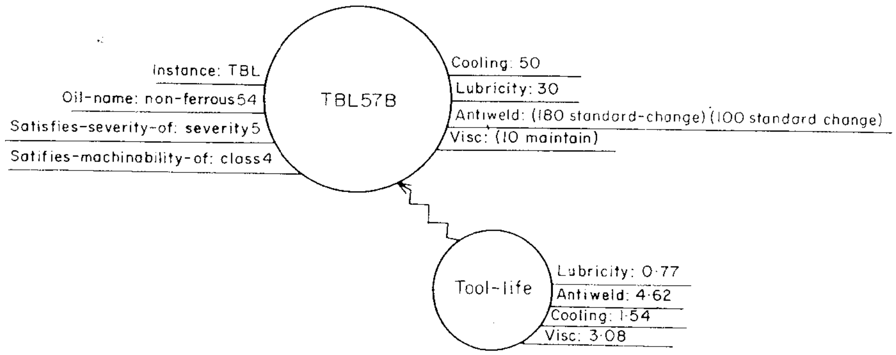

FIGURE 22. Empirical selection causal network example

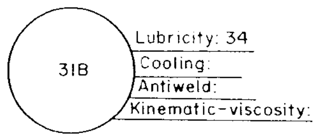

FIGURE 23. Example cutting fluid—'31B'

The utility function which measures how well an actual value of lubricity matches the optimal value is stored in the GOAL-FUNC slot of **GOAL1**—(i.e. MAINTAIN-HIGH).

With this introduction, the steps the fluid evaluator performs in calculating the lubricity rating of a given cutting fluid (e.g. '31B') relative to a given process (e.g. PROCESS-SPEC2) are as follows:

It finds the tool life goal attached to the **LUBRICITY** schema. In this case, the tool life goal is represented by the **GOAL1** schema;

It extracts the 'optimal' value of lubricity residing in the DEGREE slot of the **GOAL1** schema (i.e. 64);

It extracts the value of the LUBRICITY slot of the cutting fluid in question—31B (i.e. 34)—Figure 23, and gives this as an argument to the rating function residing in the GOAL-FUNC slot of the **GOAL1** schema (i.e. **MAINTAIN-HIGH**) along with the 'optimal' lubricity value.

The value returned from the rating function serves as a measure of how well the lubricity of cutting fluid 31B approximated the desired 'optimal' value for the process—PROCESS-SPEC2. This value is then multiplied by the sensitivity coefficient for lubricity whose determination was explained earlier, to give the overall contribution of lubricity to tool-life for the process PROCESS-SPEC2. The other properties, cooling, viscosity, and antiweld, are determined similarly.

5.5 AVERAGE SHOP RATING DETERMINATION

	Recommendation results				
	Tool-life	Lubricity	Cooling	Antiweld	Viscosity
Sensitivities:	10.01	0.77	1.54	4.62	3.08
31A	9.35	0.70	1.15	4.57*	2.92
31B	9.10	0.58	1.17	4.44*	2.91
41M	8.95	0.72	1.28*	4.11	2.84
41D	8.20	0.73*	0.62	4.07	2.77*
41E	failed NO-CHLORINE				

The fluid evaluator evaluates each high level property relative to a given process. It then combines these individual values into one composite value. The calculation of the tool life value will be used as an example.

An estimate of how a given cutting fluid will affect the tool replacement costs of a given job sho[p] depends upon three factors:

how the fluid affects the life of each individual too[l] (*tool-life*);[21]

the percentage of jobs in the shop that ar[e] performed with that tool (*process pct*);[22]

how much that tool costs (or the tool lif[e] importance—*importance*).[23]

The equation combining these quantities into a[n] overall rating of the fluid is:

Shop rating

$$= \frac{\sum_{i=1}^{n} \text{process } pct_i \times \text{tool life}_i \times \text{importance}_i}{\sum_{i=1}^{n} \text{importance}_i}$$

The equation gives an average rating for the fluid weighing the fluid's performance for n processes in th[e] shop.

6. GREASE system evaluation

Three sets of experiments were performed to evaluat[e] the performance of GREASE subsequent to th[e] primary phases of knowledge acquisition and enhan[ce]ment. The experiment sets were designed to test th[e] major capabilities of GREASE and included:

cutting fluid recommendations for single sho[p] processes without diagnostic symptoms and atypica[l] operating conditions;

recommendations for single shop processes with [a] variety of symptoms and atypical operatin[g] conditions;

recommendations for shop containing multipl[e] processes for which a single cutting fluid is desired

Each experiment set tests progressively mor[e] features of GREASE and its success depends upo[n] the success of the prior sets of experiments.

The experiment sets include test cases selecte[d] both from actual customer shops and hypothetica[l] shops. Individual experiments consisted of performin[g] the same shop recommendations independently b[y] GREASE, a cutting fluid expert, and a cutting flui[d] salesman. The salesman represents an experience[d] non-expert, who will be the primary user of th[e] GREASE system.

[21] The calculation of how the fluid affects the life of a tool wa[s] explained in the previous section.

[22] The percentage of jobs in the shop performed by a given tool is [a] simple calculation based upon the shop representation.

[23] How much a given tool costs cannot be calculated by GREASE However, the customer is prompted for an estimate of the relativ[e] cost of a tool on a scale from 1 to 10. The default is 1.

Analysis of experimental results consisted of comparing the recommendations from each source to answer two primary questions:

how well does GREASE agree with the experts and experienced salesmen?

when GREASE returns a high rating for a fluid, does that really mean that in an expert's opinion it will do a good job?

To answer these questions two tests were performed for each recommendation:

comparison of the top three choices from the different sources followed by an explanation of their differences;

expert rating of the fluids that GREASE gave a high value on its rating scale. The expert ratings are good, satisfactory, or poor, and are based solely on expected performance.

6.1 EXPERIMENTAL DATASETS

GREASE was tested with several actual field cases. These cases represent a wide range of problems for GREASE and test all features for a variety of commonly encountered customer shops, excluding diagnosis and compensation for atypical operating conditions. Several hypothetical cases were analysed by GREASE and the cutting fluid expert to test these features.

The field cases fall into two categories:

single process recommendations;

multiple process recommendations.

The field cases included a range of machining operations from the highest severity (i.e. broaching) to the lowest severity (i.e. grinding) that GREASE can consider and a wide range of materials including members from all the ferrous groups and selected non-ferrous materials.

The multiple process cases tested the ability for GREASE to properly average a shop for which a single cutting fluid is desired. This average depends on the ability of GREASE to properly recommend a single process shop and, in selected cases, to average both ferrous and non-ferrous recommendations.

A description of the field cases is in Figure 24.

A series of hypothetical test cases were devised to test the ability of GREASE to properly recommend fluids if diagnostics or atypical operating conditions were specified in the customer's shop. Actual field cases were unavailable, since it is currently not common procedure to collect this information. These cases are described in Section 6.2. For the diagnostic cases, the current fluid '41B' was assumed. In order to observe the primary affect of the diagnostic or atypical condition upon the recommendations and to minimize scatter, the material machined and machine operation were held constant in these cases.

Case#	Materials machined	Operations
1	100% GROUP1—1144, LEADED STEELS	100% REAMING
2	40% GROUP2—12-L-14, 12-L-15, 1144 40% GROUP3—6150, 8620 20% GROUP4—52100, 4140	100% AUTOMATICS-MULTIPLE-SPINDLE
3	10% GROUP1—416, 4130 20% GROUP3—410, 420 70% GROUP4—440C STAINLESS	100% MILLING
4	20% GROUP1—1117, 1212, LEADED STEELS 20% GROUP3—303 STAINLESS 60% GROUP6A—BRASS, BRONZE	100% AUTOMATICS-MULTIPLE-SPINDLE
5	20% GROUP1—LEADED STEELS 10% GROUP2—CAST IRON 70% GROUP4—HIGH CARBON—1040, M2	80% AUTOMATICS-MULTIPLE-SPINDLE 20% DRILLING
6	100% GROUP4—52100, 440C	100% FORM-GRINDING
7	80% GROUP3—203 STAINLESS 20% BRASS	50% AUTOMATICS-MULTIPLE-SPINDLE 50% TAPPING
8	100% GROUP1—LEADED STEELS, 1144	100% GUN DRILLING-bore: 0.218in.
9	100% GROUP4—M2 STEEL	100% CENTERLESS-GRINDING
10	100% GROUP3—4047 LEADED	100% EXTERNAL BROACHING
11	100% GROUP4—4047 (NOT-LEADED)	100% DRILLING
12	100% GROUP4—M1, M7 STEELS	100% FORM-GRINDING
13	100% BRASS	100% AUTOMATICS-MULTIPLE-SPINDLE

FIGURE 24. Actual field test cases used to evaluate GREASE

Case#	Material machined	Operation	Diagnostic or Operation conditioned
14	100% GROUP1	MULT-SPINDLE AUTO.	(control case)
15	100% GROUP1	MULT-SPINDLE AUTO.	HIGH-SPEED
16	100% GROUP1	MULT-SPINDLE AUTO.	HIGH-SPEED, DEEP-CUT
17	100% GROUP1	MULT-SPINDLE AUTO.	HIGH-SPEED, DEEP-CUT, HIGH-FEED
18	100% GROUP1	MULT-SPINDLE AUTO.	LONG CHIPS
19	100% GROUP1	MULT-SPINDLE AUTO.	LONG CHIPS, LARGE BUE
20	100% GROUP1	MULT-SPINDLE AUTO.	LONG CHIPS, LARGE BUE, TOOL-SEIZURE

FIGURE 25. Diagnostic and atypical conditions test cases

6.2 EXPERIMENTAL EVALUATION

Two sets of experimental results were collected for each experiment. The first dataset represented the best three cutting fluid choices for the customer's shop.[24] The results are compared to determine how good an agreement exists between GREASE, the expert, and the salesman in terms of a fluid which meets the customer's needs. The intent is to determine if GREASE performs worse, on a par, or better than its human counterparts.

The second dataset is a ranking by the expert of the high rated fluids recommended by GREASE for the customer shop. The results give an indication of how well a high-performance level fluid selected by GREASE might actually perform in a customer's shop.

6.2.1 Best three choices

The three best cutting fluid choices independently determined by the expert salesman, and GREASE were correlated in table form. The selected fluids are indicated, followed by a 'rating' value which is an estimate by either the expert or salesman of how well this fluid would perform in the customer's shop.[25] The 'rating' for the GREASE selections was determined by the expert. The 'rating' values are G—good, S—satisfactory and P—poor. It is important to note that three good choices might not always be possible in the cutting fluid product line, in which case rating values less than good are indicated. The 'rank' indicates where the expert- and salesman- selected fluid falls in the GREASE recommendation table. The value 't.l.', for 'tool-life', indicates the perfor mance level of the specified fluid as determined b GREASE. The 'score' of a particular choice is th percentage of the maximum rating possible (i.e., whe all choices receive a 'good' rating).

Lastly, cutting fluids which are currently used in customer's shop are indicated with an asterisk in th first fluid choice of the salesman.[26]

6.2.2 High rated fluids

For each recommendation made by GREASE, th expert rated the fluids with tool-life ratings greater o equal to 8.00 in terms of how well they would perfor in the customer's shop. The ratings were G—good S—satisfactory and P—poor.

6.3 THIRD EXPERIMENT SET—MULTIPLE PROCESSES

The third set of experiments tested GREAS for shop recommendations containing multiple pro cesses for which a single cutting fluid is desired (Detailed results for experiments 1 and 2 can be foun in Mogush *et al.*, 1986). These experiments tested th ability of GREASE to correctly combine proces ratings of cutting fluids for individual processes withi the shop into an overall shop average. In developin the shop average, GREASE assumes a linearity in th single process rating scale and uses process percent ages and process importances to weigh the average The field cases tested the ability of GREASE t average within the same material classification (e.g ferrous materials), as well as between materi classifications comprising both ferrous and non-ferrou materials.

[24] All actual test cases, except case No. 13, were restricted to non-soluble cutting fluid recommendations, since the customers' shops were unequipped to handle soluble fluids.

[25] Experimentally testing each cutting fluid selection in a physical shop is impossible. Best estimates, based upon experience and expertise of cutting fluid properties, were used for the ratings.

[26] The salesman generally recommended the current fluid as the fi choice if he was aware of its identity.

6.3.1 Analysis of best three choices

The experimental results for the best three choices are in Figures 26, 27, and 28. Analysis of the results revealed that GREASE, the cutting fluid expert, and the salesman made good recommendations as the first choice. All recommended fluids would perform well in the customer's shop. The expert also was more in agreement with GREASE than the salesman on the first choice.

GREASE was on a par with the expert or salesman in terms of its second choice. All selections would result in good or satisfactory performance in the customer's shop.

GREASE generally performed well on its third choice compared to both the expert and salesman.

	Grease			Expert				Salesman			
Case#	Fluid	Rating	t.l.	Fluid	Rating	Rank	t.l.	Fluid	Rating	Rank	t.l.
2	31A	G	9.53	31B	G	3	9.30	31B*	G	3	9.3
3	31A	G	8.70	31A	G	1	8.70	31C	G	2	8.63
4	TS991	G	9.75	41D	G	1	9.75	41D*	G	5	8.15
5	31A	G	9.23	31A	G	1	9.23	31B	G	2	8.80
7	TS991	G	9.00	45A	G	5	8.25	31B*	G	—	
Average			9.24			2.2	9.05			3.0	8.72
	100% Good				100% Good				100% Good		
Score:	100%				100%				100%		

FIGURE 26. Multiple processes—first choice

	Grease			Expert				Salesman			
Case#	Fluid	Rating	t.l.	Fluid	Rating	Rank	t.l.	Fluid	Rating	Rank	t.l.
2	TS991	G	9.32	45A	S	11	7.95	31A	G	1	9.53
3	31C	G	8.63	31B	G	3	8.20	31A	G	1	8.70
4	41M	G	9.66	45A	G	10	6.94	41E	S	9	7.90
5	31B	S	8.80	41M	S	8	8.12	31A	G	1	9.23
7	41M	S	8.52	41M	S	2	8.52	31A	G	—	
Average			8.99			6.8	7.95			3.0	8.84
	60% Good				40% Good				80% Good		
	40% Satisfactory				60% Satisfactory				20% Satisfactory		
Score:	80%				70%				90%		

FIGURE 27. Multiple processes—second choice

	Grease			Expert				Salesman			
Case#	Fluid	Rating	t.l.	Fluid	Rating	Rank	t.l.	Fluid	Rating	Rank	t.l.
2	31B	G	9.30	41M	P	6	8.91	41E	P	7	8.75
3	31B	G	8.20	41M	S	9	6.90	31B	S	3	8.20
4	#372	P	9.24	43B	S	5	8.21	41B	S	4	9.13
5	TS991	S	8.67	41D	S	5	8.53	31C	G	4	8.62
7	41D	S	8.48	41D	S	3	8.48	41E	S	6	8.20
Average			8.78			5.6	8.21			4.8	8.58
	40% Good				0% Good				20% Good		
	40% Satisfactory				80% Satisfactory				60% Satisfactory		
	20% Poor				20% Poor				20% Poor		
Score:	60%				40%				50%		

FIGURE 28. Multiple processes—third choice

	Choice	Grease Score	Grease t.l.	Expert Score	Expert t.l.	Salesman Score	Salesman t.l.
Single Process:	1	100	9.76	100	8.99	100	7.98
	2	94	9.41	78	8.45	87	7.48
	3	56	7.82	36	7.17	57	7.38
Atypical conditions:	1	100	9.66	100	9.46		
	2	87	9.28	62	8.95		
	3	87	9.17	30	8.32		
Diagnostic conditions:	1	100	9.81	100	9.77		
	2	100	9.76	87	9.31		
	3	100	9.69	87	9.68		
Multiple Processes:	1	100	9.24	100	9.05	100	8.72
	2	80	8.99	70	7.95	90	8.84
	3	60	8.78	40	8.21	50	7.58

FIGURE 29. Summary of best three choices

6.4 OBSERVATIONS

6.4.1 Best three choices

GREASE performed very well in comparison to the expert and experienced salesman, as the summary of score values for the best three choices in Figure 29 demonstrates.

Analysis of the overall summary for the three best choices along with individual test cases reveals some interesting conclusions:

- Identical fluids were not always recommended for each test that resulted in identical fluid ratings. The reason is that there are generally multiple fluids in each performance class (e.g. good, satisfactory, poor) for a particular shop.
- GREASE can often find more fluids that have good performance in a shop than either the salesman or expert. The reasons for this include:
  - GREASE considers the *entire* production line for each cutting fluid selection. The salesman and expert often consider only a group of fluids that are generally used for a particular machine operation on a specific material without considering the merits of fluids not designed for a particular application.
  - In a multiple process case where a single fluid is desired for several machining processes, GREASE attempts to rigorously calculate the relative need of each individual process, rather than using estimation.
- The first choice of GREASE, the salesman, or expert always resulted in good performance in the customer's shop.
- The second choice of GREASE always resulted in good or satisfactory performance, but the expert and salesman sometimes made a poor performance choice.
- The third choice had some poor performance selections by GREASE, the expert, and the salesman.
- There was some general disagreement between the expert and salesman.

6.4.2 High rated fluids

The analysis of recommendations with GREASE projected high tool-life ratings is to determine whether such selections will do a good job in customers' shops. Two datasets were collected for this analysis:

- tool-life values for the best three choices for GREASE, the expert, and the salesman,[27]
- for every test case, the expert rated each GREASE recommendation with tool-life values greater than 8.0 into the categories of G—good, S—satisfactory, or P—poor, which reflect expected performance in a customer's shop.

Analysis of tool-life ratings for the best three fluids revealed:

- there exists a rough correlation to tool-life values with rating 'scores' within each experiment (i.e. a lower rating 'score' results in a lower tool-life value);
- tool-life values decrease with choice number within an experiment;
- an absolute tool-life value could not be associated with a rating 'score' across all experiments including

[27] The tool-life values for the expert and salesman are GREASE projected values for their fluid selections.

recommendations by GREASE, the expert and the salesman. An example of this is seen in the multiple process case in Figure 29. The first choice for GREASE, the expert, and the salesman all receive ratings of 100% , but the tool-life ranged from 9.76 for GREASE to 7.98 for the salesman. The 7.98 value is greater than the value that GREASE reports for its third choice—7.82 which results in a rating score of only 56%.

The dataset comprising expert ratings of GREASE recommendations with tool-life ratings greater than 8.0 was compiled for all 20 test cases. Individual cases were not examined since it was the intent to determine if tool-life values by themselves could be correlated with good, satisfactory, or poor recommendations. A table was prepared relating tool-life range vs. number of observations of good, satisfactory, or poor performance, Figure 30.

Analysis of tool-life vs. performance reveals:

tool-life values greater than 9.4 always resulted in good performance;

a range of tool-life values for satisfactory without poor performance could not be determined;

a wide range of tool-life values exists for each performance classification. For example, the average tool-life value for 'good' was 9.5 with a standard deviation of 0.9.

a large overlap of performance classes exists.

Tool-life range	Good	Satisfactory	Poor
9.9–10.0	5	0	0
9.8	12	0	0
9.7	15	0	0
9.6	9	1	0
9.5	2	0	0
9.4	3	3	0
9.3	8	2	1
9.2	2	1	3
9.1	1	3	4
9.0	1	2	2
8.9	0	2	2
8.8	0	3	2
8.7	1	1	2
8.6	2	4	2
8.5	0	9	3
8.4	0	4	5
8.3	0	4	2
8.2	3	2	5
8.1	0	4	1
8.0	0	3	1

FIGURE 30. GREASE tool-life values vs. performance

6.4.3 Observations concerning GREASE evaluation by expert and salesman

Interviewing of the expert and salesman after the experiments revealed the following:

they were satisfied with their cutting fluid choices for the test cases;

they sometimes didn't consider a fluid which GREASE recommended because they either did not think of it or they wouldn't have even considered it for a particular application. Some fluids new to the product line were often overlooked. They were excited over the utility of GREASE to be able to identify potential applications of fluids prior to actual field usage;

they both generally considered GREASE to be very useful for performing cutting fluid recommendations;

there were differences in selections made by GREASE and the expert despite the fact that knowledge input of GREASE was obtained from the expert. There are two reasons for this:

GREASE is able to more rigorously calculate the effectiveness of a fluid for a given process;

GREASE was designed to make conservative predictions and give the best choice in all cases, whereas the expert, in many cases, chose a slightly poorer performing oil which would be more cost effective in terms of performance and price.

7. Comparison with previous work

The treatment problem has received little attention in the literature. This is due primarily to the diagnostic problems having simple treatments relative to the difficulty in performing the diagnosis. In the following, the treatment problem of three diagnostic systems are reviewed.

7.1 COMPARISON WITH MYCIN

In MYCIN (Shortliffe, 1976) the problems are basically broken down into a number of sub-problems (one for each micro-organism present). Solutions to these sub-problems are found and combined into a global solution. Then the global solution is tested to make sure that global constraints, e.g. age, weight, health of the patient, and allowable drug combinations, are satisfied. Searching is stopped once a satisfactory solution is found.

There are a number of important differences between these treatment problems which make the treatment problems faced by GREASE significantly more difficult.

Unlike MYCIN, GREASE has to treat more than one 'patient' at once. This means that there may not be any one ideal solution for all of the problems present.

When MYCIN presents a possible remedy, it is simply the selection of a particular drug. However, when GREASE suggests a possible remedy, there are still a multitude of ways of realizing that goal. For example, an increase in lubricity can be accomplished by any number of different combinations of increases in fatty content and increases in viscosity.

In MYCIN, there are a comparatively large number of loosely coupled solutions to each individual problem. This means that if individual problems are solved separately, the probability is high that some combination of the individual solutions will yield an acceptable global solution. On the other hand, in GREASE, the number of possible solutions to problems are small (usually involving the variation of one of four properties), and the solutions are very tightly coupled. This tight coupling occurs in two ways. First of all, a solution to one problem might call for increasing a given property, while the solution to another problem might call for decreasing the same property. This is not unlikely because of the small number of properties to manipulate. Secondly, different properties are also coupled with each other; fluids low in viscosity tend to be low in lubricity for instance. This means if the solution to one problem calls for a low viscosity and the solution to another problem calls for a high lubricity, the solutions will conflict with each other. This makes it unlikely that an "ideal" fluid which could solve all of the problems perfectly exists.

In MYCIN, there are a comparatively large number of possible global solutions (combinations of drugs). In GREASE, the number of possible global solutions is limited to the product line.[28] This, together with the tight coupling of the solutions in GREASE, makes it very unlikely that the available global solutions to GREASE problems are "ideal".

The MYCIN treatment problem is a satisfying problem whereas the GREASE treatment problem is an optimizing problem. MYCIN can divide possible solutions into exactly two categories: those that are satisfactory and those that are no satisfactory. The problem for MYCIN is to find satisfactory treatment. GREASE, as pointed ou above, is highly unlikely to be able to find a: "ideal" solution in the existing product line. It mus evaluate solutions on a continuous scale c adequacy. The criteria by which GREASE mus evaluate a fluid are, in turn, satisfied to a greater o lesser degree on a continuous scale. The probler for GREASE is to find the optimal product in th product line.

[28] If cutting fluid additives are considered, the number of possible solutions would be extended. However, selections are generally restricted to a very limited product line. Additives are considered if a new product is being designed to be added to the line.

7.2 GENAID

GENAID is a realtime, sensor-based diagnosti system for turbine generators (Osborne, 1986; Fox ε *al.*, 1983). It uses a version of MYCIN's causa network in a forward propagation mode to identif system faults. It is able to alter the causal networ automatically based on identified degradations c sensors. Selection of treatment is quite simple. Once problem is identified, a pre-defined repair procedur exists for correcting it.

7.3 PIES

The PIES system (Pan and Tenenbaum, 1986 diagnoses problems in semi-conductor fabricatio processes by analyzing parametric test data. I employs a multi-level causal structure to represent th relationships between

- parametric measurements,
- physical silicon structure,
- fabrication process, and
- malfunctions in: fabrication equipment, sourc materials, environment and human operation.

At each level, cases describing failure modes defin the failure and the strength of its causal relations t failures at the same and other levels.

Both diagnosis, i.e., identifying the cause of th failure mode, and treatment are viewed as one an the same. This is due to either an assumed one-to-on relationship between root causes and treatment or th ability to experiment with alternative treatments t identify which corrects the problem.

7.4 TEST

TEST (Kahn, 1987) is a troubleshooting shell whic has been applied to a number of applications includin automobile troubleshooting. The troubleshootin concept differs from the diagnosis approach taken in PIES, where all the symptoms are known (provided by the test equipment) prior to diagnosis. In PIES, diagnosis reduces to the propagation of support through a causal network. In the case of troubleshooting, complete symptom information is not available *a priori,* but must be gathered incrementally and at low cost; planning the sequence of tasks to perform becomes important.

TEST uses a causal network centered around failure modes with embedded expertise to guide the reasoner in selecting which causal path to pursue. As with the PIES system, the identification of the root cause of a failure is sufficient to identify a unique treatment given the one-to-one relationship between cause and discrete repairable function.

8. Conclusion

Many decision problems are composed of two parts: diagnosis and treatment. Much of the work in expert systems has focused on the use of heuristic classification to perform diagnosis. The treatment of the diagnosed problem is either canned or takes a divide-and-conquer approach assuming independence among solutions. In the cutting fluid selection domain, the cross-product of materials, operations, and requirements is so large that a causal network cannot be constructed to relate them to the available fluids. The task was further complicated by the need to use a single fluid for many different operations.

Because of the lack of a complete causal network, expert knowledge was used to construct a theoretical ideal for each machining operation/material pair. The properties of this "starting point" were then heuristically modified based on shop and diagnostic constraints. The "heuristic optimum" was then used as an evaluation function to rate the fluids in the product line. Consequently, the treatment process can be viewed as a combination of heuristic and analytic techniques where an evaluation function (i.e., distance metric) is heuristically determined before searching for a fluid. Whether the search process is selective or synthetic, the same evaluation may be used.

In closing, one could pose the question of whether GREASE is an 'AI' system. It is certainly the case that knowledge in the form of expertise is used to construct the evaluation function. Secondly, the AI paradigm of symbolic knowledge representation has played an important role in modeling the domain. On the other hand, the actual rating of fluids is basically a weighted distance metric; nothing fancy is happening there. The important lesson is not whether GREASE is an AI system, but whether the use of AI in conjunction with more conventional techniques can solve the problem. The answer is: yes it can.

References

Aikins, J. S., Kunz, J. C., Shortliffe, E. H. and Fallat, R. J. 1983. PUFF: an expert system for interpretation of pulmonary function data. *Computers and Biomedical Research* **16,** 199–208.

American Society for Metals 1968. *Metals Handbook,* Vol. 3, 8th ed.

Clancey, W. J. 1984. Classification problem solving. In: *Proceedings of the National Conference on Artificial Intelligence,* pp. 49–55. Los Angeles: Morgan Kaufmann.

Duda, R. O., Hart, P. E., Barrett, P., Gashnig, J., Konolige, K., Reboh, R. and Slocum, J. 1978. Development of the PROSPECTOR Consultant System for Mineral Exploration. Final Report for SRI Projects 5821 and 6415. Technical Report, Artificial Intelligence Center, SRI International.

Fox, M. S., Lowenfeld, S. and Kleinosky, P. 1983. Techniques for sensor-bases diagnosis, In: *Proceedings of the Eighth International Joint Conference on Artificial Intelligence,* pp 158–163. Los Angeles: Morgan Kaufmann.

Gulf Oil 1981*a*. *Selection of Cutting and Grinding Fluids, Module X—Lubricants and Lubrication.*

Gulf Oil 1981*b*. *Metal Machining with Cutting Fluids.*

Gulf R&D 1982. *Gulfcut Selection Guide for Machining and Grinding Operations.*

Kahn, G. S. 1987. From application shell to knowledge acquisition system. In: *Proceedings of the International Joint Conference on Artificial Intelligence,* pp. 355–359. Los Angeles: Morgan Kaufmann.

Machinability Data Center 1972. *Machining Data Handbook,* 2nd ed. Cincinnati: Metaut Research Assoc. Inc.

Mogush, J. E., Carrega, D., Spirtes, P. and Fox, M. S. 1986. Treatment as Heuritic Construction: A Case Study in Cutting Fluid Selection. Report CMU-RI-TR-86-15, Robotics Institute, Carnegie-Mellon University, Pittsburgh.

Osborne, R. L., 1986. Online, artificial intelligence-based turbine generator diagnostics. *AI Magazine,* 97–103.

Pan, J. Y-C. and Tenenbaum, J. M. 1986. PIES: An engineer's do-it-yourself knowledge system for interpretation of parametric test data. *AI Magazine,* 62–69.

Shortliffe, E. H. 1976. *Computer-Based Medical Consultations: MYCIN.* New York: American Elsevier.

Versonder, G. T. *et al.* 1983. ACE: An expert system for telephone cable maintenance. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence,* pp. 116–121. Reading: Addison-Wesley.

James E. Mogush is a Senior Project Engineer for Pittsburgh Applied Research Corp. (PARC), a subsidiary of the University of Pittsburgh. He holds BS and MS degrees from the University of Pittsburgh. His current position as coordinator of computer engineering involves the design of MIS and real-time process monitoring and control systems for

manufacturing and business-related applications. Before joining PARC in 1986, he was a Senior Research Engineer at Gulf Research Corp., responsible for pioneering advanced computer applications in the areas of artificial intelligence, process control, and MIS.

Dominique J. Carrega graduated from the Ecole Centrale de Paris in 1983. He studied Computer Applications to Architecture and City Planning at the Architecture Department of Carnegie-Mellon University before joining the Robotics Institute, where he participated in the development of the GREASE project. Returning to France, he worked as a consultant for the Advanced Technologies Division of General de Services Informatiques and developed various expert systems. He now acts as a senior consultant for industrial artificial intelligence applications, and sales contact for Carnegie Group products in western Europe.

Peter Spirtes received a PhD in History and Philosophy of Science from the University of Pittsburgh in 1981, and an MS in Computer Science from the University of Pittsburgh in 1983. After working at the Intelligent Systems Laboratory at Carnegie-Mellon University for several years, he became an assistant professor in the Philosophy Department at Carnegie-Mellon. He specializes in the philosophy of the social sciences. He is currently developing TETRAD, an intelligent system for inferring causal structure from correlation data. He is the co-author of *Discovering Causal Structure: Artificial Intelligence, Philosophy of Science, and Causal Modeling.*

Mark S. Fox received his BSc in Computer Science from the University of Toronto in 1975 and his PhD in Computer Science from Carnegie-Mellon University in 1983. In 1979 he joined the Robotics Institute of Carnegie-Mellon University as a Research Scientist. In 1980 he started and was appointed Director of the Intelligent Systems Laboratory. In 1984, he co-founded Carnegie Group Inc., a software company which specializes in knowledge-based systems for solving engineering and manufacturing problems. Carnegie-Mellon University appointed him Associate Professor of Computer Sciences and Robotics in 1987. In 1988 he was appointed Director for the new Center for Integrated Manufacturing Decision Systems, which is one of the largest centers in the U.S. for research in intelligent systems to solve engineering and manufacturing problems. Dr Fox pioneered the application of artificial intelligence to factory planning and scheduling problems, project management, and material design. He was the designer of PDS/GENAID, a steam turbine generator diagnostic system, which was a recipient of the IR100, and was the creator of SRL from which Knowledge Craft, a commercial knowledge engineering tool, was derived. Research interests include knowledge representation, constraint directed reasoning and applications of artificial intelligence to engineering and manufacturing problems. Dr. Fox has published over 50 papers, and is a member of AAAI, ACM, IEEE, SME, CSCSI and TIMS.