INTEGRATION OF MULTIPLE KNOWLEDGE SOURCES IN ALADIN, AN ALLOY DESIGN SYSTEM

M. D. Rychener,[1] M. L. Farinacci,[2]
I. Hulthage,[1] and M. S. Fox[1]
[1]Intelligent Systems Laboratory, Robotics Institute,
Carnegie-Mellon University, Pittsburgh, PA 15213
[2]ALCOA Laboratories, ALCOA Center, PA 15069

ABSTRACT

ALADIN* is a knowledge-based system that aids metallurgists in the design of new aluminum alloys. Alloy design is characterized by creativity, intuition and conceptual reasoning. The application of artificial intelligence to this domain poses a number of challenges, including: how to focus the search, how to deal with subproblem interactions, how to integrate multiple, incomplete design models and how to represent complex, metallurgical structure knowledge. In this paper, our approach to dealing with these problems is described.

1. INTRODUCTION

ALADIN (ALuminum Alloy Design INventor) is a knowledge-based system that aids metallurgists in the design of new aluminum alloys. The alloy design task produces a material composition and thermal-mechanical processing (TMP) plan whose resulting alloy satisfies a set of criteria, e.g., Ultimate Tensile Strength. The system can be operated in several modes. As a decision support system, it accepts alloy property targets as input and suggests alloying additives, processing methods or microstructural features to meet the targets. As a design assistant, it can evaluate designs supplied by a metallurgist, or provide information that is useful for design from a knowledge bank. As a knowledge bank, it provides information to supplement the usual sources such as books, journals, databases and specialized consultants.

Alloy design in an industrial setting involves teams of experts, each of whom is a highly-trained specialist in a different technical area. The primary application objective of ALADIN is to systematize and preserve the expertise of such teams, as an expert system. There is some hope that by fusing together multiple sources of knowledge from different experts, a system will be developed that exceeds the capabilities of individual experts. At the same time the expertise can be applied more widely to design problems. We also hope with such a system to shorten the design cycle, which is often on the order of five years, from specification of properties until commercial production begins.

Alloy design raises a number of issues as an AI problem. *First*, the search space is combinatorially complex due to the number and amount of elements that may participate in the composition and the number of alternative processing plans**. The knowledge available to guide the search is primarily heuristic, gained over many years of experimentation, coupled with some metallurgical models. As a result, *there exist multiple partial models of alloys* which relate:

- composition to alloy properties,
- thermal-mechanical processing to alloy properties, and
- micro-structure to alloy properties.

This raises two questions for AI: what is the appropriate architecture for the explicit representation and utilization of multiple, parallel theories, and how is search to be focused in this architecture?

A *second* issue is the degree to which design decisions are dependent. Each change in composition or process alters a number of properties. This level of dependence results in a level of interaction among sub-problems which exceeds that experienced in the planning literature, and is not amenable to simple constraint propagation techniques due to the size and complexity of the search space.

Issue *three* is the result of issues one and two. The complexity of the search places a tremendous burden on how to focus attention in complex solution spaces.

Lastly, issue *four* is concerned with representation. Knowledge of the relationship between alloy structure and its resultant properties is at best semi-formal. Much of it is composed of diagrams of 3D structure and a natural language description. Quantitative models rarely exist. The problem lies in representing spatial information in which structural variations are significant.

The rest of this paper describes the alloy design problem in more detail. This is followed by a description of the ALADIN problem-solving architecture. Then there is a discussion of knowledge representation, multi-model reasoning, and focus of attention.

2. ALLOY DESIGN REASONING

An alloy design problem begins with the specification of constraints on the physical properties of the material to be created. The objective of the designer is to identify element additions with percent levels and processing methods that will result in an alloy with the desired characteristics. The line of reasoning that designers use is similar to the generate-and-test model. The designer selects a known material that has properties similar to the design targets or other interesting features. The designer then alters the properties of the known material by making changes to the composition and processing methods. The effects of these changes on the various physical properties are estimated, and discrepancies are identified to be corrected in a later iteration.

In order to select fabrication variables that improve the

*This research has been supported by the Aluminum Company of America.

**The MOLGEN system [14] focused primarily on process planning.

properties, the designer may consider known cause and effect relations, such as:

- *IF* Mg is added *THEN* the strength will increase
- *IF* the aging temperature is increased beyond the peak level *THEN* the strength will decrease

Often, however, these relationships are not available or cannot be generalized sufficiently. In that case, the designer may construct a model of the microstructure that will produce the required properties. The microstructure can be defined to be the configuration in three-dimensional space of all types of non-equilibrium defects in an idealized phase. These defects include voids, cracks, particles and irregularities in the atomic planes. They are visible when the material is magnified several hundred times with a microscope. The geometric, mechanical and chemical properties of the microstructural elements, as well as their spatial distributions and interrelationships, have a major influence on the macroscopic properties of the material. The microstructure is often described in abstract, conceptual terms and is rarely characterized numerically. However, these concepts provide a powerful guide for the search process since they constrain composition and processing decisions. For example, if meta-stable precipitates are required, then the percentage of additives must be constrained below the solubility limit, certain heat treatment processes must be applied, and aging times and temperatures must be constrained within certain numerical ranges.

While the human design approach can generally be characterized with the generate-and-test model, a more detailed study of metallurgical reasoning reveals complexities and deviations from the idealized artificial intelligence-based models. To some extent, knowledge is applied in an opportunistic fashion. When relationships or procedures are identified that can make some progress in solving the problem, then they may be applied. However, there are many regularities to the search process. Furthermore, the strategies that designers use to select classes of knowledge to be applied varies among individuals. For example, in the selection of the baseline alloy to begin the search, some designers like to work with commercial alloys and others prefer experimental alloys produced in a very controlled environment. Still others like to begin with a commercially pure material and design from basic principles. When searching for alternatives to meet target properties, some designers construct a complete model of the microstructure that will meet all properties and then they identify composition and processing options. Many designers prefer to think about one property at a time, identifying a partial structure characterization and implementation plan that will meet one property before moving to the next. Still other designers prefer to avoid microstructure reasoning whenever possible by using direct relationships between decision variables and design targets. All designers occasionally check their partial plans by estimating the primary and secondary effects of fabrication decisions on structure and properties. However, the frequency of this activity and the level of sophistication of the estimation models varies among designers.

3. ALADIN ARCHITECTURE

ALADIN is a multi-spatial reasoning architecture akin to a blackboard model [3, 6]. It is composed of five spaces:

1. **Property Space:** The multi-dimensional space of all alloy properties.
2. **Structure Space:** The space of all alloy microstructures
3. **Composition Space**: This is an space where each dimension represents a different alloying element (e.g., Cu, Mg).
4. **Process Space**: The space of all thermo-mechanical alloy manufacturing processes.
5. **Meta Space**: This is the focus of attention planning space which directs all processing. The meta space holds knowledge about the design process and control strategies. Planning and search takes place in this space in that goals and goal trees are built for subsequent execution.

Each space can be viewed as a separate blackboard system with its own search space, hypotheses and abstraction levels. Activity is generated on different planes and levels in a way similar to Stefik's MOLGEN system [14]. ALADIN's planes are: **Meta** or strategic plane, which plans for the design process itself, establishing sequencing, priorities, etc.; **Structure** planning plane, which formulates targets at the phase and microstructure level, in order to realize the desired macro-properties; **Implementation** plane, encompassing chemical composition and thermal and mechanical processing subplanes.

We treat the alloy design problem as a planning problem because the final alloy design is a sequence of steps to be taken in a production plant in order to produce the alloy. The design plan is only partly ordered since the time ordering of some steps is unimportant. The planning process in ALADIN utilizes the existence of the microstructure model. The alloy design therefore typically starts in the structure space with decisions on microstructural features that imply desirable properties. These decisions are thereafter implemented in composition and process space. Overall, the search is organized according to three principles that have proven successful in past AI systems:

- **Meta Planning**
- **Least Commitment:** meaning that values within hypotheses are expressed as ranges of values that are kept as broad as possible,
- **Multiple levels:** under which plans are developed first at an abstract level, and then gradually made more precise.

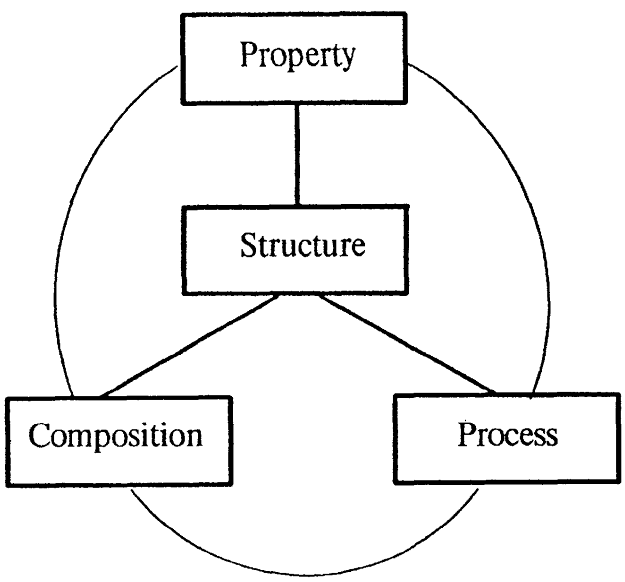

Figure 3-1: Spaces of Domain knowledge

properties, the designer may consider known cause and effect relations, such as:

- *IF* Mg is added *THEN* the strength will increase
- *IF* the aging temperature is increased beyond the peak level *THEN* the strength will decrease

Often, however, these relationships are not available or cannot be generalized sufficiently. In that case, the designer may construct a model of the microstructure that will produce the required properties. The microstructure can be defined to be the configuration in three-dimensional space of all types of non-equilibrium defects in an idealized phase. These defects include voids, cracks, particles and irregularities in the atomic planes. They are visible when the material is magnified several hundred times with a microscope. The geometric, mechanical and chemical properties of the microstructural elements, as well as their spatial distributions and interrelationships, have a major influence on the macroscopic properties of the material. The microstructure is often described in abstract, conceptual terms and is rarely characterized numerically. However, these concepts provide a powerful guide for the search process since they constrain composition and processing decisions. For example, if meta-stable precipitates are required, then the percentage of additives must be constrained below the solubility limit, certain heat treatment processes must be applied, and aging times and temperatures must be constrained within certain numerical ranges.

While the human design approach can generally be characterized with the generate-and-test model, a more detailed study of metallurgical reasoning reveals complexities and deviations from the idealized artificial intelligence-based models. To some extent, knowledge is applied in an opportunistic fashion. When relationships or procedures are identified that can make some progress in solving the problem, then they may be applied. However, there are many regularities to the search process. Furthermore, the strategies that designers use to select classes of knowledge to be applied varies among individuals. For example, in the selection of the baseline alloy to begin the search, some designers like to work with commercial alloys and others prefer experimental alloys produced in a very controlled environment. Still others like to begin with a commercially pure material and design from basic principles. When searching for alternatives to meet target properties, some designers construct a complete model of the microstructure that will meet all properties and then they identify composition and processing options. Many designers prefer to think about one property at a time, identifying a partial structure characterization and implementation plan that will meet one property before moving to the next. Still other designers prefer to avoid microstructure reasoning whenever possible by using direct relationships between decision variables and design targets. All designers occasionally check their partial plans by estimating the primary and secondary effects of fabrication decisions on structure and properties. However, the frequency of this activity and the level of sophistication of the estimation models varies among designers.

3. ALADIN ARCHITECTURE

ALADIN is a multi-spatial reasoning architecture akin to a blackboard model [3, 6] . It is composed of five spaces:

1. **Property Space:** The multi-dimensional space of all alloy properties.
2. **Structure Space:** The space of all alloy microstructures
3. **Composition Space**: This is an space where each dimension represents a different alloying element (e.g., Cu, Mg).
4. **Process Space**: The space of all thermo-mechanical alloy manufacturing processes.
5. **Meta Space**: This is the focus of attention planning space which directs all processing. The meta space holds knowledge about the design process and control strategies. Planning and search takes place in this space in that goals and goal trees are built for subsequent execution.

Each space can be viewed as a separate blackboard system with its own search space, hypotheses and abstraction levels. Activity is generated on different planes and levels in a way similar to Stefik's MOLGEN system [14]. ALADIN's planes are: **Meta** or strategic plane, which plans for the design process itself, establishing sequencing, priorities, etc.; **Structure** planning plane, which formulates targets at the phase and microstructure level, in order to realize the desired macro-properties; **Implementation** plane, encompassing chemical composition and thermal and mechanical processing subplanes.

We treat the alloy design problem as a planning problem because the final alloy design is a sequence of steps to be taken in a production plant in order to produce the alloy. The design plan is only partly ordered since the time ordering of some steps is unimportant. The planning process in ALADIN utilizes the existence of the microstructure model. The alloy design therefore typically starts in the structure space with decisions on microstructural features that imply desirable properties. These decisions are thereafter implemented in composition and process space. Overall, the search is organized according to three principles that have proven successful in past AI systems:

- **Meta Planning**
- **Least Commitment:** meaning that values within hypotheses are expressed as ranges of values that are kept as broad as possible,
- **Multiple levels:** under which plans are developed first at an abstract level, and then gradually made more precise.

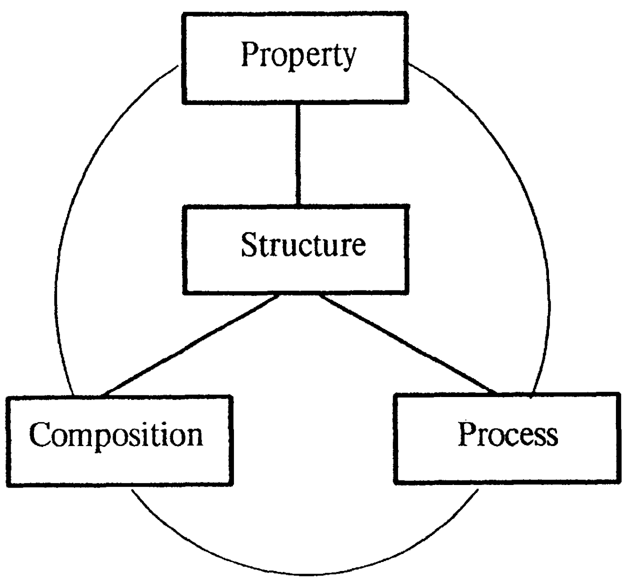

Figure 3-1: Spaces of Domain knowledge

The general trend of execution is to start generating a plan in the Meta plane, and to complete the alloy design within the processing plane. However, it will always be necessary to jump back and forth between spaces and levels and to backtrack.

The qualitative and quantitative levels of the Structure, Composition and Processing spaces are activated as appropriate, to generate hypotheses that specify design variables in their own range of expertise. Hypotheses generated on other planes and levels constrain and guide the search for new hypotheses in many ways. An existing qualitative hypothesis obviously suggests the generation of a quantitative hypothesis. Certain microstructure elements can be produced by compositional additives, while others are produced by specific processes with the composition restricting the choices available. The final product of the design process is a plan in the composition and process spaces. More details on the ALADIN architecture are available in [11].

4. KNOWLEDGE REPRESENTATION

ALADIN utilizes three forms of knowledge representations:

1. **Declarative** knowledge base of alloys, properties, products, processes, and metallurgical structure concepts;
2. **Production Rules** in the form of IF-THEN rules of many types: control of search among competing hypotheses, empirical associations of causes and effects, rankings and preference orderings, processing of user commands, decisions about when to call upon knowledge in other forms, and others;
3. **Algorithms** knowledge expressed as functions: detailed physical, chemical, thermo-mechanical, statistical, etc. calculations.

4.1. Declarative Knowledge

The ALADIN system contains representations for metallurgical charts, alloys, physical properties, compositions and processing methods [8]. Each of these classes of knowledge admit a relatively simple representation using well known ideas about schemata (frames) and inheritance. The representation of microstructure presents some interesting problems and is discussed in more detail here.

Microstructure is the configuration in three-dimensional space of all types of non-equilibrium defects [7] in an ideal phase. Metallurgical research has shown that many microstructural features have important consequences for macroscopic properties. The objective of the microstructure representation in ALADIN is to classify and quantify the microstructure of alloys in order to facilitate the formulation of rules that relate the microstructure to the macroscopic properties of alloys.

Although much of the heuristic knowledge about alloy design involves the microstructure, it is usually poorly represented. Metallurgists have attempted to describe microstructural features systematically [7] and there is also a field called quantitative metallography that describes quantitative information about the three-dimensional microstructure of alloys [15]. In practice, neither of these approaches is commonly used. Instead, metallurgists rely on visual inspection of micrographs, which are pictures of metal surfaces taken through a microscope. Information is communicated with these pictures and through a verbal explanation of their essential features.

In order to represent microstructure data and rules it was necessary to develop a symbolic representation of alloy microstructure. The two main features of an alloy microstructure are the grains and the grain boundaries, and are described by an enumeration of the types of grains and grain boundaries present. Each of these microstructural elements are in turn described by any available information such as size, distribution, etc., and by its relations to other microstructural elements such as precipitates, dislocations, etc. This representation allows important facts to be expressed even if quantitative data is unavailable, an important example being the presence of precipitates on the grain boundaries. It is interesting to note that most of the expert reasoning about microstructure deals with qualitative facts, with quantitative information typically not available.

4.2. Procedural Knowledge

Most of the procedural knowledge is encoded in OPS5 production rules [5, 1] in well-known ways. However, some procedures are best represented as algorithms, for whose coding we have chosen Common Lisp [13]. Especially important is the fact that alloy design requires the simultaneous use of both qualitative or symbolic reasoning and the application of suitable mathematical models. The subject of coupling symbolic and numeric methods is of general interest. Accordingly, Kitzmiller and Kowalik [10] point out that in order to solve many problems in business, science and engineering, both insight and precision are needed.

ALADIN currently contains the following types of mathematical routines:

- Regressions, in order to interpolate and extrapolate from known alloy properties to those of new alloys;
- Models of structure-insensitive properties, such as density;
- Solutions of systems of multi-dimensional constraints;
- Retrieval of constraints from phase diagrams.

ALADIN couples qualitative and quantitative reasoning in several ways. The design is made at two levels, first on a qualitative and second on a quantitative level. Examples of design decisions that are made first are what alloying elements to add and whether the alloy should be artificially aged or not. These decisions are followed on the quantitative level by a determination of how much of each alloying element should be added and at what temperature aging should take place.

The ALADIN system attempts to couple symbolic and numeric computation deeply by not treating algorithms as black boxes. A calculation is typically broken down into calculations of the various quantities involved, and the exact course of a computation is determined dynamically at the time of execution through the selection of methods to determine all the quantities needed to obtain the final result. These selections are based on heuristic knowledge that estimates the relative advantage and accuracy of the choices and by the availability of data [9].

5. MULTIPLE DESIGN MODELS

It is a feature of the alloy design domain that several partly independent models of alloys are used. The simplest model of alloys deals only with the relationship between chemical composition and alloy properties. From the point of view of modern metallurgy only a few structure-independent properties like density and modulus can be described in this way. However, empirical knowledge does exist about other properties, eg. Beryllium causes embrittlement in Aluminum. Quantitative comparisons can also be made between alloys of varying composition, everything else being equal. This yields some useful quantitative knowledge about properties through regression.

A more complete model includes the relationship between thermo-mechanical processes and properties. Since only composition and process descriptions are needed to manufacture an alloy, it could be assumed that no other models are needed to design alloys. As a matter of fact, historically many alloys have been designed with composition and process models only. The progress of research in metallurgy is giving new insights in the relationship between the microstructure of alloys and their physical properties. The deepest understanding of alloy design therefore involves models of microstructure effects on properties and models of composition and processing effects on microstructure.

The microstructure decisions serve as an abstract plan that cuts down the number of alternatives in the composition and process spaces. In this way the role of the microstructure has both similarities and differences with abstract planning as described by Sacerdoti [12]. The main differences are:

- Microstructure concepts are distinct from composition and process concepts, not merely a less detailed description.
- The microstructure plan is not a part of the final design in the sense that an alloy can be manufactured with composition and process information only.
- The microstructure domain is predefined by metallurgical expertise, not defined during implementation or execution of the ALADIN system.

These differences introduce a number of differences from a MOLGEN-like system:

- Instead of one hierarchy of plans there are three; Structure, Composition and Process, each of which has abstraction levels.
- Since structure decisions don't necessarily always have the highest criticality (as defined by [12]), opportunistic search is important.
- The effect of abstract hypotheses is more complex because decisions in the structure space cut the search by constraining the choice of both composition and process hypotheses. The existence of more than one level in each space also introduces new types of interactions.

Ideally, the models taking microstructure into account should be sufficient for all design decisions, but in reality they are incomplete. As a result, empirical models that relate composition and processes directly to properties have to be used. Utilizing several design models introduces another important deviation from standard abstract hierarchical planning: One or more levels of the Structure space can be bypassed during hypothesis generation or property evaluation.

It is the combined use of the five design models plus a set of global control strategies for dealing with multiple models that enables ALADIN to design an alloy.

6. DESIGN STRATEGY PLANNING AND FOCUS OF ATTENTION

ALADIN has a model of alloy design strategies that is encoded in OPS5 rules and associated with the meta space. This space is used to guide and control the search for solutions. The strength of this strategic model comes from the partition of the detailed metallurgical knowledge into knowledge sources. Facts, rules and procedures are each associated with a knowledge source that is characterized by a context, a goal, a space and a level. Rules and procedures can be applied only if the corresponding goal and context is active.

The design strategy model guides the search by building goals. Several types of information are included:

- The status of the search,
- The history of the solution process,
- Constraints on strategic alternatives, and
- The effectiveness of various strategic alternatives.

The status of the search is characterized by the constraints, hypotheses and estimates that have been created and indicate what problems remain to be dealt with. These schemata have the following definitions:

- **Hypothesis.** Partial description and commitment regarding the alloy that is designed to meet the targets.
- **Constraint.** The design target, and therefore a condition to be met and a criterion for selecting hypotheses.
- **Estimate.** Prediction of the effect of fabricating an alloy according to the components of the current hypothesis; the effects will show up as characteristic properties and microstructure.

The history of the solution process is retained in the goals. ALADIN has a rather elaborate set of rules for managing goals in a general way. Each goal has a status, a symbol from a fixed set of possibilities, which are in turn understood and managed by general goal rules. The outcome of work on a goal is propagated according to its final status and the logical (e.g., AND and OR) and sequential relations (nextgoal) that the goal has to other goals.

Constraints on control alternatives are easily represented in rule form. Some examples are:

- If numerical decisions regarding composition and process have not yet been made, then quantitative evaluation models can not be applied;
- If decisions have not yet been made regarding what processing steps to use, then it makes no sense to reason about temperatures and rates.

Finally, the system has a notion of what strategies will have the greatest impact on the search, based on heuristic knowledge obtained from the metallurgists. Rules include:

- If it is possible to reason about microstructure, composition or process, then microstructure reasoning is preferred;
- If many fabrication alternatives have been identified to meet a single target, then use simple heuristics to evaluate each and prune the search.

Due to the complex interdependence of design decisions on an alloy's final properties, simple concepts of goal protection are inappropriate. Instead a combination of least commitment and *over-compensatory* planning is utilized. This means to over-compensate when achieving a goal. In particular, if a certain tensile strength is required, the planning system sets even higher goals to achieve at this point in the search knowing that later decisions may result in a reduction of this property. This approach works because the property goals are values on a continuum.

ALADIN begins in the meta space and frequently returns there for new direction. When the meta space is activated, strategy rules identify activities that are reasonable and create top level goals in memory, with context, space, and level information. Often, several alternative strategies are possible at any point in the search, and the user is offered a menu of possibilities. The system recommends the strategy that is felt to be most effective. After the user makes a selection, the meta rules expand the goals by creating more detailed subgoals. These goal trees constitute a plan for how to accomplish the requested activities. Control then returns to the domain spaces, which process the goals. Control remains in the domain spaces until the success or failure of each goal is determined. At that point, control returns again to the meta space. Iteration between meta and domain space continues until the ALADIN problem solving process is complete.

With the meta space, numerous design strategies, obtained from different people, are integrated into a single system. As a result, ALADIN can develop several solutions to a single problem by applying different approaches. The flexible user control allows the metallurgist to experiment with different strategies. The designer may, in fact, explore solutions arising out of the application of hybrid strategies that are not usually integrated into a single problem.

7. SYSTEM PERFORMANCE AND RESULTS

ALADIN runs on a Symbolics Lisp Machine within the Knowledge Craft [2] environment at a speed that is comfortable for interaction with expert alloy designers. A typical design run takes about an hour, and involves considerable interaction with the user, whose choices influence the quality of the outcome. Its development is at the mature, advanced-prototype stage, where it can begin to assist in the design process. We must point out, though, that its knowledge is presently focused on narrow areas of alloy design, with expertise on only three additives, two microstructural aspects and five design properties. We are dealing in depth only with ternary alloys. But these restrictions are by our own choice, so that we can go into depth and train the system on the selective areas of greatest import to our expert informants and sponsors. Within these restrictions lie a number of commercially important alloys, whose rediscovery and refinement by ALADIN will be a major milestone.

Performance measures to date are strictly anecdotal. Our experts work with the system in the interactive mode described earlier. Three milestones have been reached:

1. The representation of structural knowledge is considered by the experts to be an advance over what was available previously.
2. The experts have made the transition from being sceptics to believing the system is of value to their work.
3. The system is beginning to produce non-trivial results that are of interest to designers, and that would require too much tedious work to generate manually. These include partial designs on several spaces and levels.

Though two years have passed since the commencement of the project, we continue to work with the experts to refine and extend the voluminous knowledge and data not yet added to the system. More details about the current state of user acceptance and future plans for technology transfer are supplied in [4].

ALADIN is primarily an application of existing artificial intelligence ideas to an advanced, difficult problem domain. Alloy design is thought to require a high degree of creativity and intuition. However, we have found that generate-and-test, abstract planning, decomposition and rule-based heuristic reasoning can reproduce a significant portion of the reasoning used by human designers on prototype cases. Furthermore, the attempt to build a knowledge-based system has helped alloy designers to systematize their knowledge and characterize interrelationships, particularly in the area of microstructural representation.

9. ACKNOWLEDGMENTS

We are grateful to all our expert metallurgist informants from Alcoa, especially Marek Przystupa and Warren Hunt.

10. REFERENCES

1. Brownston, L., Farrell, R., Kant, E., Martin, N.. *Programming Expert Systems in OPS5: An Introduction to Rule-Based Programming.* Addison-Wesley, Reading, MA, 1985.

2. Carnegie Group, Inc. Knowledge Craft, Version 3.0. Pittsburgh, PA, October, 1985.

3. Erman, L. D., Hayes-Roth, F., Lesser, V. R. and Reddy D. R. "The Hearsay-II speech-understanding system: integrating knowledge to resolve uncertainty". *Computing Surveys 12,* 2 (June 1980), 214-253.

4. Farinacci, M. L., Fox, M. S., Hulthage, I. and Rychener, M. D. The development of Aladin, an expert system for Aluminum alloy design. Third International Conference on Advanced Information Technology, Amsterdam, The Netherlands, 1986. Sponsored by Gottlieb Duttweiler Institute, Zurich, Switzerland, November, 1985; also Tech. Rpt. CMU-RI-TR-86-5.

5. Forgy, C. L. OPS5 User's Manual. CMU-CS-81-135, Carnegie-Mellon University, Dept. of Computer Science, July, 1981.

6. Hayes-Roth, B. "A blackboard architecture for control". *Artificial Intelligence 26* (1985), 251-321.

7. Hornbogen, E. "On the microstructure of alloys". *Acta Metall. 32,* 5 (1984), 615.

8. Hulthage, I., Farinacci, M. L., Fox, M. S., Przystupa, M., and Rychener, M. D. The Metallurgical Database of Aladin. Carnegie-Mellon University, Intelligent Systems Laboratory, Robotics Institute, 1986. In preparation.

9. Hulthage, I., Rychener, M. D., Fox, M. S. and Farinacci, M. L. The use of quantitative databases in Aladin, an alloy design system. Coupling Symbolic and Numerical Computing in Expert Systems, Amsterdam, The Netherlands, 1986. Presented at a Workshop in Bellvue, WA, August, 1985; Also Tech. Rpt. CMU-RI-TR-85-19.

10. Kitzmiller, C. T. and Kowalik, J. S. Symbolic and numerical computing in knowledge-based systems. Coupling Symbolic and Numerical Computing in Expert Systems, Amsterdam, The Netherlands, 1986.

11. Rychener, M. D., Farinacci, M. L., Hulthage, I. and Fox, M. S. Integrating multiple knowledge sources in ALADIN, an alloy design system. Carnegie-Mellon University, Intelligent Systems Laboratory, Robotics Institute, 1986. In Preparation; Long version of AAAI-86 paper.

12. Sacerdoti, E. D. "Planning in a hierarchy of abstraction spaces". *Artificial Intelligence 5* (1974), 115-135.

13. Steele, G. L.. *Common Lisp, the Language.* Digital Press, Burlington, MA, 1984.

14. Stefik, M. J. "Planning with constraints (MOLGEN: part 1); Planning and meta-planning (MOLGEN: part 2)". *Artificial Intelligence 16* (1981), 111-170.

15. Underwood, E. E.. *Quantitative Stereology.* Addison-Wesley, Reading, MA, 1970.