

The Development of ALADIN, an Expert System for Aluminum Alloy Design

**Martha L. Farinacci*, Mark S. Fox,
Ingemar Hulthage, and Michael D. Rychener**

CMU-RI-TR-86-5

**Intelligent Systems Laboratory
The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213**

***Process Control and Computer Technology Division
Alcoa Laboratories
Alcoa Center, Pennsylvania 15069**

January 1986

Copyright © 1986 Carnegie-Mellon University

This research was funded by the Aluminum Company of America.

This report has appeared in *Proceedings, Third International Conference on Advanced Information Technology*. Gottlieb Duttweiler Institute, Zürich, Switzerland, November 1985.

Table of Contents

1. Introduction	1
2. Assessment of Artificial Intelligence for Alloy Design	1
3. Alloy Design Reasoning	3
4. Case Study of ALADIN	4
4.1. First Iteration	5
4.2. Second Iteration	8
5. Conclusions	11
6. References	11

List of Figures

Figure 3-1: Alloy design graph	4
Figure 4-1: First iteration flow of control	7

Abstract

ALADIN is an expert system that aids metallurgists in the design of new aluminum alloys. The project began in January 1984, and two major iterations of the system design have been completed. The current system is a hybrid of several approaches and artificial intelligence techniques. Declarative representations are used for metallurgical data and concepts. The basic design procedure is encoded in a rule-based, with the potential for flexible control and participation by the user. The system is based on the hypothesize-and-test model, and a constraint-based search is used to find alternative designs. Strategies are included for resolving the interactions and trade-offs among conflicting property targets. Metallurgical relationships cover a broad range of symbolic and numerical types of reasoning, which are integrated in the system. In this article, an overview of the project is given, including discussions of knowledge acquisition techniques, design decisions and implementation.

1. Introduction

ALADIN (ALuminum Alloy Design INventor) is an expert system that aids metallurgists in the design of new aluminum alloys. The system can be operated in several modes. As a decision support system, it will accept alloy property targets as input and suggest 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.

ALADIN was developed through the combined efforts of three organizations: The Intelligent Systems Lab of the Robotics Institute of Carnegie-Mellon University, the Artificial Intelligence group at Alcoa Laboratories, and the Alloy Technology Division of Alcoa Laboratories.

Design and development of the system began in January, 1984 and continues at the time of this writing. While it is difficult to evaluate a system that is under development, it is hoped that ALADIN will provide many important benefits when completed. The amount of knowledge required to successfully develop new materials is so great that individuals often must supplement their private knowledge with information from books, journals and specialized consultants. ALADIN, when used as a knowledge bank, will provide another source of valued information. 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. The development of new alloys now requires testing many alternatives before one adequate solution is found. An attempt is being made to develop a more complete understanding of the behavior of alloys through the development of quantitative models. These models will, in turn, enable development to proceed with fewer tests. ALADIN will serve as a collection point for these models and also the older empirical methods. Hence ALADIN will aid in the development of new materials at a reduced cost and time.

In this report, a case study of the ALADIN project to date is given.

2. Assessment of Artificial Intelligence for Alloy Design

Computers have long been used to solve engineering and science problems in the metals industries. For example, finite element models of numerous manufacturing processes, including rolling, extrusion, forming, forging, casting and smelting, have been developed [9]. With these models, the effects of processing on the shape of the material, internal stresses, temperatures, and currents (in the case of smelting) are predicted. Simulation methods are usually used to study processing on a larger scale. Like finite element methods, these models are predictive although they often deal with multiple processing steps. They can determine the effects of material handling practices, plant layout, equipment design and operating practices on product characteristics and equipment performance. Another major application of computing in the metals industry is process or numerical control. Tool adjustments are made in order to control the shape and quality of the product being produced. Often the control models are based, in part, on the equations underlying the predictive models mentioned earlier [2].

The classical applications of computing just described share several important features:

- The problems are primarily numerical.

- The processes are well understood quantitatively, and equations are available to describe effects.
- Well defined algorithms are available to solve the equations and make the appropriate predictions.

Engineering problems with these qualities are usually solved with models implemented in a procedural language such as FORTRAN.

Alloy design, in contrast to these problems, is characterized by the following features:

- Quantitative models for calculating the properties of a proposed design are often not available, or insufficient data prevent their use.
- Models that are available contain a high degree of uncertainty.
- The design process, as practiced by people, often involves reasoning about images appearing under a microscope, abstract concepts, interpretations and heuristic rules of thumb based on experience. The reasoning is highly symbolic as opposed to numeric.
- The alloy design process, as practiced by people, appears to be similar to the diagnostic problems discussed in artificial intelligence literature. When given a new set of alloy specifications, designers identify an existing alloy, determine in what way the existing alloy fails to meet specifications, and try to find a cause for the discrepancy. Alloy design problems are more complex than diagnostic problems, however, since a great deal of reasoning is required to construct a solution even after the cause of the problem is identified.
- No fixed algorithm seems to be used in design. Instead, metallurgists seem to search through the space of design variables with guidance from heuristic rules and models. The flow of reasoning depends on what knowledge is available and how reliable that knowledge is at the time the design problem is specified.

It was because of these features that an artificial intelligence approach was used in the alloy design problem.

Although alloy design has a few features that make it a natural artificial intelligence application, there are other features that make automation difficult. For example, the solution of a single problem takes weeks or often months. This is because the reasoning is complex, many options must be explored, and vital missing information must be searched for. Because of this, it was recognized that knowledge acquisition would be difficult. Furthermore, design involves reasoning about time. Processing of the metal involves changes to the structure of the material which influence later production steps. In other words, the time and sequence in which changes are made is important. It is known that planning problems of this sort are difficult to solve. However, it was felt that the benefits of an alloy design system would outweigh the risks involved in its development.

3. Alloy Design Reasoning

Before beginning a technical discussion of the ALADIN system, an overview of the reasoning used in alloy design will be given. This overview does not require any knowledge of metallurgy. The introduction of a few basic scientific concepts will make later descriptions of the system clearer.

The alloy designer usually begins his problem with a set of constraints on the physical properties of the material he is to make. The constraints are often one-sided, as in: strength must be higher than x and density must be lower than y . Sometimes the constraints involve properties that are not numerical such as machinability or surface appearance. Specifications usually involve many properties and correspond to a material that does not yet exist. This is because the design problem comes from a customer or the corporation who would like something that exceeds the capabilities of everything on the market. The design problem may be overconstrained - there may be no material that can meet the specifications given. It is difficult to know in advance which problems are impossible to solve. The designer is expected to do the best he can and come as close as possible to meeting the targets.

Once property constraints are specified, designers usually select some baseline alloy to begin their search for a solution. Design will then involve finding alterations to the baseline that change the characteristics of the material in the direction of the targets. There are several different strategies that people use for selecting a baseline. Some look for a commercial alloy that comes as close as possible to meeting the targets, or the experimental alloy from the most recent iteration in the current design problem. Still others begin with a simple, commercially pure aluminum alloy and design from basic principles.

After the baseline is selected, designers look for changes that can be made to the material in order to improve the properties. The things that can be changed directly are:

- COMPOSITION - the elements added and the amounts
- PROCESS - the fabrication steps used, their sequence, and such specifications as method, temperature or time

There are rules that indicate the effects of these design choices on properties. Some examples are:

- *IF* an element with low atomic number is added *THEN* density will decrease
- *IF* Mg is added *THEN* strength will increase

However, the most powerful rules involve reasoning about the microstructure of the alloy. Some examples of these rules are:

- *IF* The aging process is done for a long time *THEN* equilibrium precipitates will form
- *IF* equilibrium precipitates are present *THEN* they are usually incoherent
- *IF* incoherent precipitates are present *THEN* they may form on the grain boundary
- *IF* precipitates are on the grain boundary and strength is medium or high *THEN* elongation and fracture-toughness are low

Metallurgists often illustrate the knowledge they work with in a graph (figure [3-1]). Nodes correspond to the four classes of knowledge that characterize an alloy: the composition, processing, structure and properties. Arrows indicate the relationships that are used so heavily in design. The rules just mentioned are examples of these relationships. For instance, the first rule: "IF an element with low atomic number is added THEN density will decrease", corresponds to the arrow from composition to property.




Figure 3-1: Alloy design graph

Designers seem to apply rules in an opportunistic fashion. Whenever rules are identified that will make some progress in solving the problem, those rules are applied. There are regularities to the search process, however. Rules that involve reasoning about structure are given a preference over rules that do not. Also, some property targets are viewed as being more important than others, and rules that deal with important targets are used first.

Because metallurgists are asked to develop materials with properties outside of the range of existing alloys, they must reason about designs which have never been tested and whose behavior is not known with certainty. Designers try to fill this gap in existing knowledge with general models, speculation, extrapolation of known trends, and analogy with existing materials. Even after applying these methods, however, it is usually impossible to determine exactly what composition and processing specification will meet the targets. The designer identifies a family of alloys that is likely to meet the targets, and makes these alloys in the laboratory. Sometimes, one member in the family will have the desired characteristics. Often, all tests will fail to meet targets exactly, but analysis of the data results in more accurate rules that can form the basis for a better set of experiments in the next iteration.

4. Case Study of ALADIN

ALADIN was designed and developed using a rapid prototype approach. During the first year, 1984, metallurgists were interviewed, the alloy design methods were characterized, a system was designed and the program written. At the end of 1984, a demonstration of the system was given to the metallurgists and an assessment of the project was made. This assessment involved technical design decisions as well as staffing levels and the approach used for knowledge acquisition. Several changes were made to the design and to the organization of the project during 1985. The design changes are now being incorporated into the software, and it is expected that the second iteration will be completed by the end of 1985.

4.1. First Iteration

During the first iteration, many alloy designers were involved in the knowledge acquisition process. During the first two months, five people were interviewed who had done work on a variety of different alloy systems. An attempt was made to find the common approaches used by all designers. By focusing on these, it was hoped that the system could be developed in a general way and could handle many different design strategies. After the initial interviews, a single prototype problem was identified and short term goals were established to develop a system based on the prototype. The problem selected was the first iteration in the design of aluminum lithium alloys. These alloys are now under development at Alcoa and are targeted for future aerospace applications. This choice was made for several technical and nontechnical reasons. The aluminum lithium alloys are important to Alcoa, so management support was expected to be high. Furthermore, many experts were available and anxious to help on the project. The reasoning used on the aluminum lithium project covered a broad range of empirical and theoretical ideas, so a system based on the prototype could be extended to other problems easily. Finally, since the company had only recently begun research on this alloy system, it was speculated that there were many design alternatives, not yet explored by human designers, that the ALADIN system could study. After the prototype problem was selected, a few more experts were interviewed, and a team of five metallurgists was finally selected for more detailed knowledge acquisition.

During the summer and autumn of 1984, the first iteration ALADIN system was designed. Because of tight deadlines for the implementation, several simplifications were made.

The solution method was modelled as a variant of generate & test [7]. This involved iteration of three steps in a fixed sequence:

- Hypothesis-Generation - find actions to move the alloy in the direction of targets
- Hypothesis-Selection - select the best option
- Hypothesis-Evaluation - predict the properties of the alloy and identify deviations from target

Metallurgy rules defining the relationships between composition, processing, structure and properties provided the operators for generation and evaluation. In cases where operators were not available, regression analysis was used to identify useful trends from the alloy data base [8]. The available generation operators often dealt with property targets independently. So, the design problem was decomposed into subproblems - one for each target. The effects of subproblem interaction were not addressed, although it was known that properties were interdependent. The search space was limited to composition choices. Few microstructure and no process rules were built into the first system. Because of this decision, the system contained almost no advanced scientific knowledge. But the developers were able to test and evaluate the use of various problem solving paradigms on this problem before spending too much time learning difficult metallurgical concepts. A hierarchical planning approach was used: an abstract plan containing the types of elements to add was developed before the details of additive percents were derived.

A representation was developed for a database of alloy information and other metallurgical knowledge. The representation was based on the concept of frames. The database included known

alloys with information about alloying and impurity elements, physical properties, product forms and applications. Alloy families and series were also stored and linked to the individual alloys. Inheritance across links was used for default reasoning. The alloy database was used to determine standard practices and to identify useful relationships between composition and properties. Representations were also developed for important metallurgical tables and diagrams. For example, phase diagrams indicate the types of equilibrium phases that are present in an alloy for a given composition and temperature. These diagrams are used extensively by alloy designers when selecting alloying additions and processing conditions. Regions of the diagram and important characteristics of the associated phases were represented in the database. Relations between regions were defined and corresponded to transformations between phases.

Two types of data elements were created to hold information about the alloy design problem - the constraint and the hypothesis. Property targets specified by the user were represented as constraints. Although in real problems, early decisions constrain choices that are available later in the search, no attempt was made to represent these dynamic constraints in the first iteration. A representation was also developed for what was called a hypothesis tree. All design choices were posted on this tree, which was developed in a depth first manner. This hypothesis tree was slightly different from the standard search graphs used for artificial intelligence problem solvers [10]. Separate trees were developed for the abstract and the detailed plans. Links were created between nodes in the two trees to indicate dependencies between the two levels of decisions.

Three types of control elements were created for the system - contexts, goals and tasks. A context corresponded to a major phase of the design problem [1]. Six contexts were identified for the first iteration system - problem-definition, search-setup, hypothesis-generation, hypothesis-selection, hypothesis-evaluation and search-termination. Contexts could have a status of active or suspended, and they were activated in a fairly rigid control sequence, as indicated in the flow chart in figure [4-1]. All metallurgy rules checked for the presence of some active context as the first precondition. Hence, the context was used to decompose the system into major subsystems. Goals were used to represent actions to be performed or already attempted [1]. Goals could be arranged into fairly general and/or/all trees and were built dynamically as the domain rules identified required subtasks. In other words, any goal could have several subgoals, and the activation of subgoals would depend on the type of the supergoal. If the supergoal was an "or" type, subgoals would be activated until one succeeded, and the success of any one subgoal would result in the success of the supergoal. If the supergoal was an "and" type, subgoals would be activated until one failed, and the success of all subgoals would be required for the success of the supergoal. Finally, if the supergoal was an "all" type, all subgoals would be activated, regardless of success or failure. In this case, the supergoal would always succeed. Sequence restrictions could be imposed on goals that shared the same supergoal, although these specifications were not required. Several status values were allowed for each goal, including posted (an action that should be performed in the future), active (an action that should be performed now), success (an action that was completed successfully) or failure (an action that was attempted but not completed). All metallurgy rules checked for the presence of an active goal as the second precondition, so the goals allowed for a finer decomposition of the rule set into knowledge sources. A general set of goal management rules was written to update status values throughout the tree whenever the status of a leaf goal node was changed by the domain rules. Finally, task elements were used to represent simple actions to be performed. Unlike goals, tasks could not be linked to form a complex tree structure.

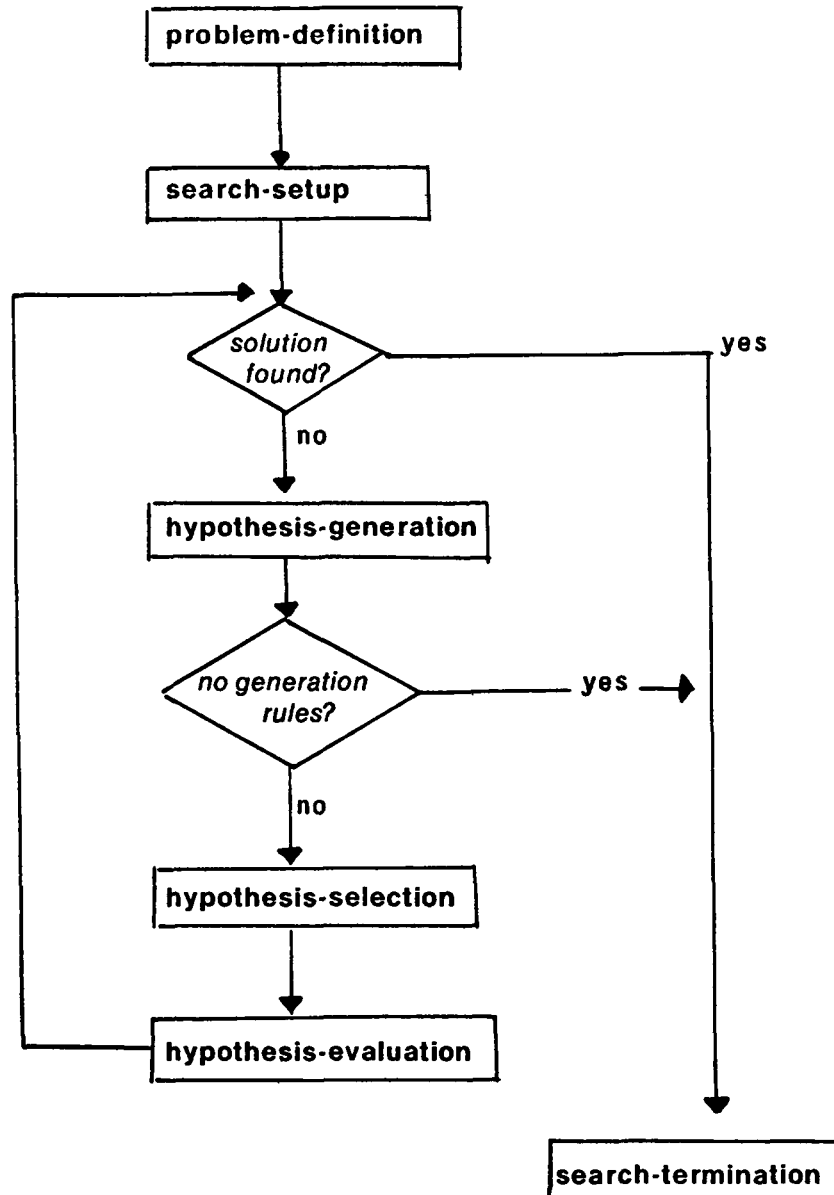


Figure 4-1: First iteration flow of control

Several conventions and standards were established for the use of control and data elements. Some examples were:

- at any time, only one context and goal should be active
- hypotheses must designate at what level of detail they reside in the hierarchical plan

Of course, when developing a system, it is always difficult to guarantee that conventions are followed, and failure to follow the practices can lead to subtle bugs that are hard to find. To make the debugging task easier, a set of diagnostic checks was formulated. These rules checked for violations from the standards and reported the errors to the developers.

During the last few months of 1984, the first iteration ALADIN system was implemented. Three languages were integrated to form the system. OPS5 [6], a forward chaining production system, was used for the overall search control, the representations of goals, hypotheses, and constraints and the metallurgy rules. CRL [4], a frame-based knowledge representation system was used to store long term declarative knowledge about alloys, their composition and properties, and also metallurgical tables and diagrams. FRANZ LISP [5] was used for the interface between OPS5 and CRL, for the user interface and also for some numerical procedures. The system was implemented on a Vax 750 running under the Unix operating system.

During the first few months of 1985, an assessment of the ALADIN project was made. Most of the earlier decisions were still found to be correct ones. For example, after a year of knowledge acquisition, it was still felt that an artificial intelligence-based approach was a reasonable one to use on the alloy design problem. Similarly, the choice of languages was still felt to be a good one. Even some of the design decisions were still felt to be appropriate. The goal management system, the decomposition of the rules base with contexts and goals, the generate and test approach, the idea of using some problem decomposition according to targets, and the use of hierarchical planning were retained for the second iteration. However, several problems were also identified. The fixed sequence of context activations was unable to respond to problems such as lack of information or failure of goals. The system had to be extended to deal with decisions about structure and process. The first iteration representation of the hypothesis tree was unable to record all of the important details about structure and process decisions. The representation of constraints also had to be generalized to deal with restrictions on composition, process and structure - decisions that could be proposed, evaluated and then withdrawn during the problem solving process. Finally, some problems were identified in the organization of the team and the approach used in knowledge acquisition. Gathering information from five metallurgists was extremely confusing and it was difficult to organize the knowledge in a coherent fashion. The primary method used to teach metallurgists about ALADIN system capabilities was to run demonstrations of the program. However, this proved to be confusing as well since at any time, most planned capabilities were not implemented and also, the user interface was not well developed.

4.2. Second Iteration

The second iteration of the ALADIN system began after the assessment phase. A conversion was made to a newer version of the CRL language and Common LISP [11], running on a VAX VMS system. Then all steps were repeated - knowledge acquisition, system design and implementation. However, the approach was changed based on the first year's experiences.

Major changes were made to the organization and staffing of the project. The aluminum lithium prototype problem was replaced by an even more specific problem that was called the training case. This case consisted of four experiments that were identified and tested during the first iteration of the aluminum lithium project. The ALADIN project goal for 1985 was to characterize and model the reasoning that was used to select those four experiments. The alloy design staff was reduced to two people, who were each allowed more time to spend on the project and given more responsibilities.

One expert had a great deal of experience with the aluminum lithium training case. He was given the responsibility of helping to prepare a description of the reasoning used on the problem. One

computer scientist assisted him with this assignment. The metallurgist tried to describe the sequence in which ideas were considered and rejected and to characterize the rules that were used to make the decisions. The computer scientist tried to reformulate this reasoning in terms of the ALADIN design concepts. In other words, the flow of metallurgical ideas and the justifications for decisions were classified as hypothesis evaluation, hypothesis generation using search operators, backtracking, etc. New contexts and goals were identified during this process, and strategies for context switching were formulated. A written report was prepared in which the reasoning for the training case was described in detail. The report was partitioned into sections - approximately half of the sections contained descriptions using metallurgical language; the other half contained formulations of the same ideas using computing concepts. During each meeting, the report was reviewed and corrections were made before work continued. After six iterations, the report was reasonably complete.

Another expert was chosen for his experience in working with computers and for his research interest in developing general models of metallurgical structure/property relationships. His main responsibilities were to help design the knowledge representation system for microstructure and processing and to help incorporate his models into the system. Another computer scientist worked closely with him to identify ways that the models could be smoothly integrated into the existing system. Written reports were again used to ensure that ideas had been understood correctly. This expert was also asked to use the system and to comment on its functionality and accuracy. Because he had more experience in using computers than many of his colleagues, it was expected that this expert would be more tolerant of the simple user interface.

The design of the system was changed in four major ways. A meta space was designed that provided a more rational sequencing of contexts and goals. The hypothesis representation was generalized to deal with decisions about processing and microstructure. A least-commitment decision-making strategy, using multidimensional constraints, was incorporated into the system. Finally, several provisions were made to deal with interactions among subproblems.

The meta space was responsible for controlling the activation of knowledge sources. It developed plans for how to solve the alloy design problem by building context elements and high level goal trees. Once the meta space activated a context and goal, control was transferred to the appropriate domain knowledge source, which could build subgoals, develop the hypothesis tree or update the status of its goals to success or failure. The meta space could then build new context and goal elements based on the status of previous goals and the reasons for their success or failure. The meta level rules were formulated to reproduce the design strategy that was used in the aluminum lithium training case.

The hypothesis representation was also generalized so that decisions about processing and microstructure could be made as well as those about composition. This extension required that new links be invented to indicate the dependencies between decision nodes. Some of these dependencies were cause and effect relationships. For example, in order to meet a property target, the program would determine the required microstructure features and post those decisions as hypotheses. During a later iteration, the program would seek composition or processing options that were likely to produce the desired microstructure features. Another type of dependency came from multiple decisions that had to be selected together in order to achieve the desired effects. The first iteration hypothesis tree was essentially an "or" tree - all children of a single node were considered to be independent options. The second iteration representation allowed for "and" branches when

multiple decisions had to be grouped together and treated as a unit. The restriction of building hypotheses using only depth first search was abandoned. A breadth first search of structure space, followed by an evaluation and pruning of poor options, followed by a depth first search of composition and process spaces seemed to be a better model of the human design process.

A general representation of multidimensional constraints was developed and used to implement a least-commitment decision-making strategy. Using this approach, decisions about exact amounts of additives or processing times and temperatures were postponed. Instead, the system derived linear or piecewise linear constraints in composition/ process space to define regions where properties were acceptable. As more property targets were considered, the sizes of these regions would decrease. If the initial problem was overconstrained or incorrect decisions were made during the search, conflicting constraints would be generated and detected. This would cause the system to backtrack or relax some of the property targets. On the other hand, if the system lacked sufficient knowledge of the problem (generation operators were lacking), or initial targets were too vague, the search procedure would terminate with a fairly large region remaining. In this case, the system would identify a family of points in the region and propose them for experimentation. In summary, the multidimensional constraints served four important purposes:

- Backtracking was reduced through the use of a least-commitment decision-making strategy.
- The interactions between search variables were represented.
- The constraint regions provided a criteria for success or failure of a design proposal.
- Multiple solutions (or one optimal solution) could be found for alloy design problems that lacked sufficient specification.

During the first iteration, design problems were decomposed according to property targets. However, most property targets were known to interact. In other words, operators that improved one physical property would also degrade other properties. Without taking into account these interactions, there was a good chance that the system would loop indefinitely. Since the nature of some types of interactions were known and understood by metallurgists, this knowledge was incorporated into the hypothesis-generation operators. These operators fired on the existence of an active goal to improve some physical property. However, the operators also checked for the existence of posted goals related to interacting properties in their preconditions. In this way, preference was given to operators that did not adversely affect other property targets being pursued in the same design problem. Other more subtle property interactions were detected during the hypothesis-selection context. Hypothesis-selection was preceded by an heuristic evaluation of all target properties. Selection was based on a measure of the distance between target and estimates in multidimensional property space. Hence, the effects of hypotheses on other property targets not being directly pursued in the current iteration were at least considered.

The second iteration design has been formulated and partially implemented. It is expected that implementation will be complete by the end of 1985. The alloy designers will then aid in another evaluation of the system.

During 1986, the ALADIN system will be converted to run on a Symbolics machine at ALCOA. After that, a more suitable user interface will be developed so that more alloy design experts will have easy access to the system. The Carnegie-Mellon team will complete its work and an ALCOA project team will prepare for continued development of the system. Major changes to the architecture or design are not expected, but a great deal of work remains in order to build the knowledge bank to an acceptable level. Also, the rule system must be generalized to deal with a broader class of problems. Several new training cases will be identified, and the reasoning will be characterized using the procedures established for the first training case. The rule system will then be augmented to reason correctly for all training cases. It is expected that the Artificial Intelligence group will play a gradually decreasing role on the project and that the Alloy Technology Division will increase the level of responsibility. This is because future work will require an ever-increasing level of knowledge of metallurgy.

5. Conclusions

Several insights were gained by our work on the ALADIN project. These insights pertain to knowledge acquisition techniques, design, architecture and implementation.

With respect to knowledge acquisition, it is very important to keep the teams small and give each individual specific responsibilities. Writing a detailed script of the knowledge and reasoning used on training cases serves several purposes and will be used on other projects. The script helps to uncover errors in communication before plans proceed too far, it helps the expert to understand some of the major computing concepts used on his problem and finally, it defines requirements during system design and implementation.

Some of the design ideas contained in ALADIN can also be extended to other systems. The use of control elements to decompose a production system into smaller units and to control rule firing has been suggested and used before [1]. The goal management subsystem is particularly powerful and is flexible enough to handle a variety of situations. Similarly, the idea of using multi-dimensional linear constraints to gradually converge on the solution to a problem with a reduction in the amount of backtracking can be applied to other application areas. The diagnostic checking rules have saved a great deal of debugging time and are recommended for complex systems requiring more than one developer. Finally, the rapid prototype method is a good one to use, and it is most effective when the prototype and the initial design are kept as simple as possible.

6. References

1. Brownston, L., Farrell, R., Kant, E., Martin, N., "Programming Expert Systems in OPS5 - An Introduction to Rule-Based Programming", Addison-Wesley Publishing, 1985.
2. Bryant, G.F., "Automation of Tandem Rolling Mills", 1973, The Iron and Steel Institute.
3. Buchanan, B.G., Shortliffe, E.H., "Rule-Based Expert Systems", Addison-Wesley Publishing Co, 1984.
4. Carnegie Group Inc, "Knowledge Craft Beta Release Notes", 1985.

5. Foderaro, J.K., "The Franz LISP Manual", University of California Technical Report, 1980.
6. Forgy, C.L., "OPS5 User's Manual", Department of Computer Science Technical Report, Carnegie-Mellon University, 1981, CMU-CS-81-135.
7. Hayes-Roth, F., Waterman, D.A., Lenat, D.B., "Building Expert Systems" , Addison-Wesley Publishing Co, 1983.
8. Hulthage, I., Rychener, M., D., Fox, M., S., Farinacci, M., L., "The Use of Quantitative Databases in Aladin, an Alloy Design System", Robotics Institute Technical Report, Carnegie-Mellon University, 1985, CMU-RI-TR-85-19.
9. Kobayashi, S., "A Review of the Finite Element Method and Metal Forming Process Modelling", Journal of Applied Metalworking, vol 2, # 3, July 1982, p163-169.
10. Nilsson, N.J., "Problem Solving Methods In Artificial Intelligence", McGraw-Hill Book Company, 1971.
11. Steele, G.L., Jr., "Common LISP the Language", Digital Press, 1984.