The Architecture of ALADIN: A Knowledge-Based Approach to Alloy Design

Ingemar A.E. Hulthage, Mark S. Fox, and Michael D. Rychener
Carnegie Mellon University

Martha L. Farinacci
Alcoa

ALADIN is a knowledge-based system to aid aluminum alloy design for aerospace applications. Alloy design is a decision process in which basic composition, and thermal- and mechanical-process steps, are selected to produce an alloy displaying a set of desired characteristics (fracture toughness, for example, and resistance to stress and corrosion cracking). This is a combinatorially explosive problem, dependent upon the choice and amounts of the composition's elemental constituents and upon the selection, parameterization, and sequencing of processing steps. The quest for a new alloy is usually driven by new product requirements. Once metallurgical experts receive a requirement set for a new aluminum alloy, they search the literature for an existing alloy that satisfies them. If none exists, the experts may draw upon experiential, heuristic, and theory-based knowledge to suggest new alloys that might exhibit the desired characteristics.

ALADIN's goal is to provide a decision support environment in which expert alloy designers can efficiently explore alternative compositions and thermomechanical-process sequences. The search for a suitable alloy design can span several years, and may require many hypothesize/experiment cycles. Reducing the number of iterations (even by one) or shortening average cycle time would provide significant gains.

Alloy design can be supported in many ways. Over the last century, many alloys with varying properties have been designed. Not all experts are aware of many of these alloys. By providing an alloy database, we can help alloy designers determine if a new design is needed. Second, by identifying alloys with properties "close to the goal," we can find a starting point for extrapolative design. Third, there are some quantitative theories linking structure, composition, and property. Providing easy access to these would aid alloy designers. Finally, all alloy design experts are not created equal. Some are more "expert" than others, and their expertise covers different domains. Capturing alloy design knowledge used by various specialists — and preserving that knowledge in an accessible form — would facilitate everyone's design efforts.

0885/9000/90/0800-0056 $1.00 © 1990 IEEE

Theoretically, we should be able to determine an alloy's properties from its microstructure. Practically, however, theories are incomplete — requiring the addition of experiential knowledge to fill gaps. As a result, multiple partial models of alloy design exist, relating

- Composition to alloy properties,
- Thermomechanical processing to alloy properties,
- Microstructure to alloy properties,
- Composition to microstructure, and
- Thermomechanical processing to microstructure.

The simplest alloy models deal only with the relationship between chemical composition and alloy properties. From the viewpoint of modern metallurgy, we can accurately determine only a few structure-insensitive properties (including density and modulus) from these models. However, empirical (and less precise) knowledge exists regarding other properties. Everything else being equal, we can make quantitative comparisons (linear regression, for example) between alloys of varying composition — comparisons that yield useful quantitative knowledge about properties.

Also, somewhat more complex models describe the relationship between thermomechanical processes and properties. Since only composition and process descriptions are needed to manufacture an alloy, we could assume that no other models are needed to design alloys. Historically, many alloys have been designed with composition and process models only. Current research in metallurgy is providing new insights into relationships between microstructures and physical properties of alloys. Therefore, the deepest understanding of alloy design involves models of microstructural effects on properties along with models of composition and processing effects on microstructure.

Thus, issues of interest in the ALADIN project are (1) what is the appropriate architecture for the explicit representation and utilization of multiple alternative models, and (2) how is search in this space of multiple interacting models to be focused?

One particularly important problem is the degree to which design decisions are interdependent. Each change in composition or process alters most properties to some degree. This differs from most design disciplines, where functionality is achieved by a system of relatively independent parts. In such cases, designers can reduce design decision interdependence by providing specifications for each part.[1] Therefore, material design involves a level of interaction among goals, extending beyond the usual situations described in AI planning literature.

Another issue concerns representation. Knowledge of the relationship between alloy structure and its resultant properties is semiformal, at best — much consists of microstructure images and natural language descriptions. Quantitative models rarely exist; even when they do, they are frequently not used. We will briefly describe how ALADIN was designed to deal with these issues. Our references provide more detail.[2-6]

We will also discuss (1) the structure of ALADIN's knowledge base, which captures complex technical concepts as well as data in symbolic, frame-based terms, (2) a multilevel planning search architecture, and (3) the integration of symbolic and numerical approaches. In addition, we will illustrate ALADIN's accomplishments by examining one fragment of an interactive design session with the system.

The knowledge base

AI has been applied to numerous engineering-design fields. Although these design areas share some features, including the need to integrate heuristics with algorithmic numerical procedures, they also contain some important differences. Each engineering subfield seems to recognize the importance of representing declarative concepts, although specific needs vary. In electrical engineering, for example, representing components with their spatial and functional relationships seems to be vital.

Mechanical engineers have studied the representation of solid geometrical shapes, which is considered crucial to the successful evolution of CAD/CAM systems.[7,8] Materials science identifies the microstructure as crucial to understanding the relationship between the characteristics of materials and their composition or processing. Therefore, a powerful representation of microstructural features is vital when constructing a materials design support system.

Our representation of declarative metallurgical knowledge demonstrates that qualitative and quantitative knowledge — available to experts in various forms, including tables, diagrams, natural language, and pictures — can be given a structured representation that enables such knowledge to be utilized through well-known AI techniques. Although many AI concepts and approaches used in the representation are routine, its application to the microstructure domain appears to be a novel concept. In fact, the literature contains few attempts to define a taxonomy for describing microstructure — and no attempts whatsoever to use a taxonomy of schemata for a computerized knowledge base of microstructure information.[9]

Alcoa used a version of this knowledge base to develop Cordial, a corrosion diagnosis system.[10] Alchemist also uses a semantic network to represent plans for designing alloys, and methods that define properties and microstructure causality.[1] While our discussion focuses on aluminum, the knowledge representation framework is useful for other alloy families — and, to some extent, even for other materials.

Woods has proposed that knowledge representation approaches be judged on two features: (1) expressive

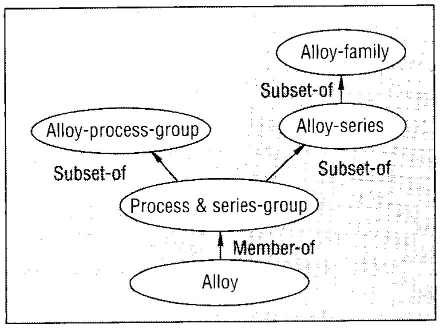

Figure 1. Alloy groups.

```
{{ 2024-T8-sheet
        INSTANCE: alloy
        MEMBER-OF: 2xxx-T8-sheet
        ADDITIVES: Cu Mg Mn
        ELONGATION: 6
        . . .
        APPLICATION: aerospace }}
```

Schema 1. A typical alloy schema.

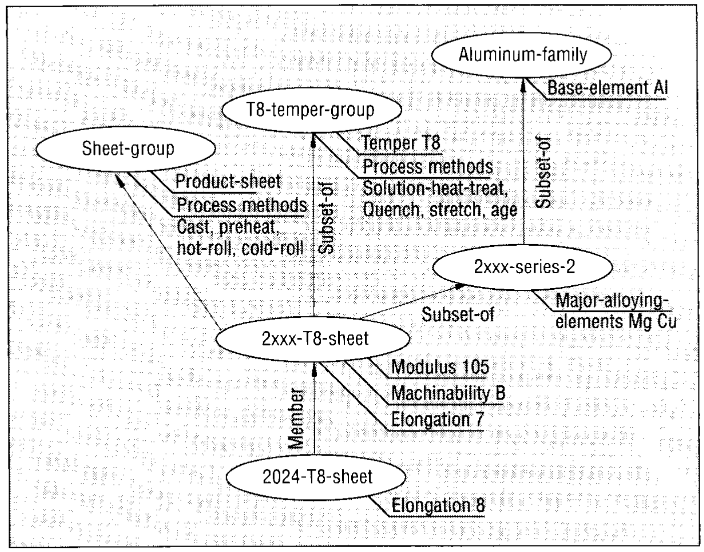

Figure 2. A representation of typical alloy properties.

adequacy, including the representation's ability to make all of the important distinctions, and to remain noncommittal about details when faced with partial knowledge; and (2) notational efficacy concerning the representation's structure and its influence on the computational efficiency of inferences, the conciseness of representation, and the ease of modification.[11]

In addition, the ALADIN representation was required to meet the following standards:

- To materials scientists, the representation should seem natural for supporting knowledge base development and maintenance by domain experts.
- The representation should be general enough to support system expansion to include non-aluminum materials.

The system organizes declarative knowledge by using schema hierarchies. Representations have a hierarchy of abstraction levels containing different degrees of detail. Knowledge Craft's facilities[12] help the system define relationships and inheritance semantics between metallurgical concepts.[13] The most commonly used relations are Is-a and Instance. The Is-a relation and some other relations define class hierarchies in which each higher level subsumes lower level classes. The Instance relation declares that a given object belongs to a class or a group; the class description serves as a prototype of instances.

The knowledge base contains information about alloys, products, applications, composition, physical properties, process methods, microstructure, and phase diagrams. The representation is general. The goal has been to create a representation for all knowledge about aluminum alloys and metallurgy relevant to the design process.

Alloy representation typifies most of the database. We will discuss it in some detail. We will also discuss microstructure, which requires a more complex representation — complexity handled largely by using Knowledge Craft's meta-information features. This enhances the representation's expressive adequacy by permitting finer optional distinctions.

The alloy hierarchy — composition, properties, and processing. When viewed from their design standpoint, alloys are interrelated and grouped together in different ways. To enable our schemata to reflect this domain organization, we have defined various formal relationships with different inheritance semantics.[13] For example, alloys are grouped together into series and families based on composition. They are also related (1) by processes that occur during fabrication — for example, heat treatment, cold rolling, and tempering, (2) by the type of application for which they are designed, and (3) by product type — for example, sheet, plate, or extrusion. Relations have been defined to reflect degrees of abstraction within the hierarchy; for instance, the relationship between a family and a prototypical member. To hypothesize and estimate, we use these relations at various points during design search — both by analogy along several different dimensions that define groups of similar alloys, and by looking for trends within these groups.

Figure 1 depicts some of this knowledge base structure: An alloy family is used to distinguish alloys by primary element (let's say, aluminum or copper). An alloy

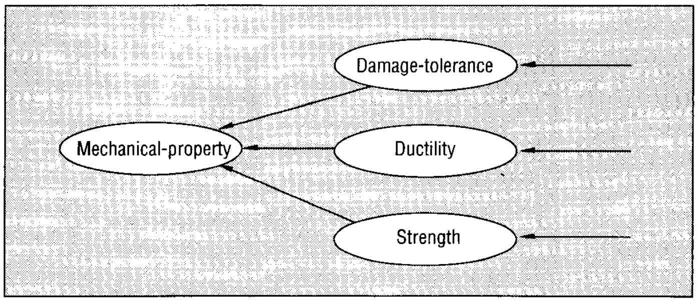

Figure 3. A mechanical-property hierarchy.

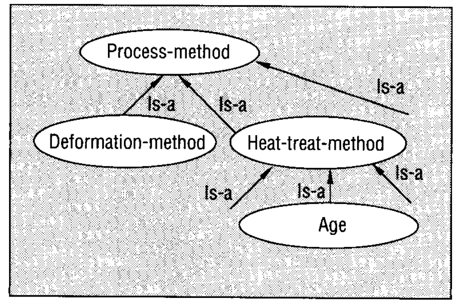

Figure 4. A process hierarchy.

series separates a family into subgroups according to secondary elements. An example of a series is those alloys containing AlMgCu. An alloy process group gathers alloys together according to processing methods (forging or casting, for example). Alloys within a series are further classified by their processing, to form a process-and-series group.

A typical alloy schema (see Figure 2 and Schema 1) will show some of the richness of representations utilized in ALADIN. In this example, the 2024-T8-sheet alloy inherits the following characteristics:

- The base element is Al, by inheritance from the aluminum family, which is an instance of an alloy family.
- The major alloying elements are Mg and Cu, by inheritance from 2xxx-series-2, which is an instance of an alloy series (Mn is considered a minor alloying element).
- The temper is T8, by inheritance from the T8 temper group, which is an instance of the alloy process group.
- The product is sheet by inheritance from the sheet group, which is an instance of the alloy process group.
- The process methods are (in order) cast, preheat, hot-roll, cold-roll, solution-heat-treat, quench, stretch, and age. The 2024-T8-sheet inherits these values from the 2xxx-T8-sheet (an instance of the process-and-series group), and the 2xxx-T8-sheet gets these values from the sheet group and the T8 temper group. Since sheet group is listed before T8 temper group in the Subset-of slot, the sheet group process methods come first.
- The modulus is 10.6 and machinability is B, by inheritance from the 2xxx-T8-sheet. The alloy could inherit the elongation value of 7 from the 2xxx-T8 sheet, but this value is overridden by the value of 6 explicitly listed with the alloy.

For use in ALADIN, we have developed a representation for more than 20 physical-property quantities. At the top classification level, properties are divided into mechanical, chemical, thermal, electrical, and miscellaneous groups. Figure 3 shows the classes of mechanical properties.

ALADIN uses the process methods classification hierarchy to make inferences about the effects of operations (on microstructure and alloy properties) since groups of methods often have similar effects. To represent operational time sequences, the system uses "before" and "after" relations. Figure 4 depicts a portion of ALADIN's process hierarchy.

Symbolic microstructure representation. ALADIN's structure knowledge falls into two categories — microstructure, and phase diagrams. We can view microstructure as the configuration of all types of non-equilibrium defects in an ideal phase.[9] Thermal and mechanical processing methods (rapid cooling and cold working) create such defects. These defects include voids, cracks, particles, and irregularities in atomic planes — features called microstructural elements, visible when the material is magnified several hundred times with a microscope. Features can range in size from a few nanometers (vacancies and Guiner-Preston zones) to micrometers (cracks and grains). Geometrical, mechanical, and chemical properties of microstructural elements — as well as their spatial distributions and interrelationships — have a major influence on the material's macroscopic properties. The objective of ALADIN's microstructure representation is to enable classification and quantification of alloy microstructures, thereby facilitating model formulations that relate the microstructure to the macroscopic properties of alloys.

Much empirical knowledge about alloy design involves the microstructure, which is difficult to represent in a useful way with standard quantitative formalisms. Metallurgists have attempted to describe microstructural features systematically[9] and quantitatively[14] but, in practice, rarely use either approach. Most expert reasoning about microstructure deals with qualitative facts. Metallurgists rely on visual inspections 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 response to this observation, we created a symbolic alloy microstructure representation that forms a crucial part of ALADIN's database.[15] Figure 5 depicts a portion of ALADIN's microstructure taxonomy. Two main features of an alloy microstructure are grains and grain

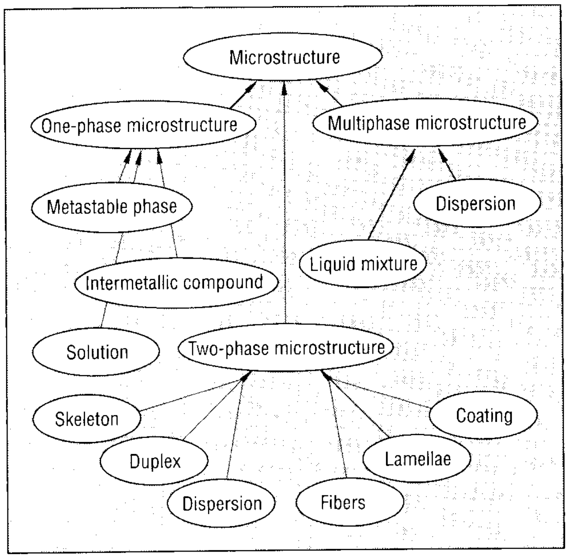

Figure 5. Classes of microstructure.

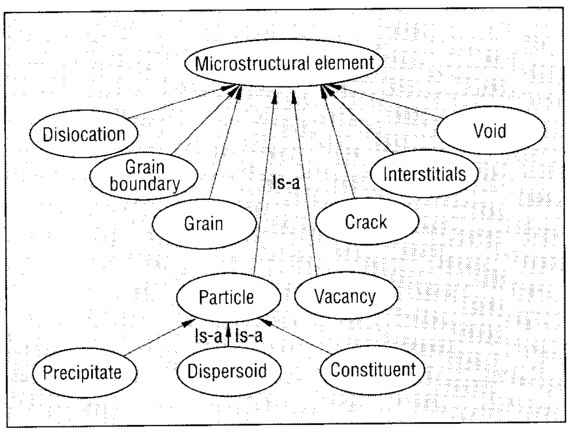

Figure 6. Types of microstructural elements.

boundaries, and these are described by an enumeration of grain types and grain boundaries present. Each microstructural element is described in turn by any available information (including size and distribution) and by relations to other microstructural elements (including precipitates and dislocations). This representation enables important facts to be expressed even if quantitative data are unavailable (the presence of precipitates on grain boundaries, for example). Figure 6 shows several such element types. Each microstructural element can be further described by its phase, size, shape, volume fraction, and distribution.

An example of microstructure representation. A typical alloy in the ALADIN database contains a microstructure description that enumerates all microstructural elements known to exist within the material. Figure 7 shows a microstructure after solution heat treatment, cold-water quenching, and peak aging at 400 degrees Fahrenheit for 48 hours.[16] Schema 2 shows the corresponding ALADIN representation of the alloy, and Schema 3 shows the microstructure.

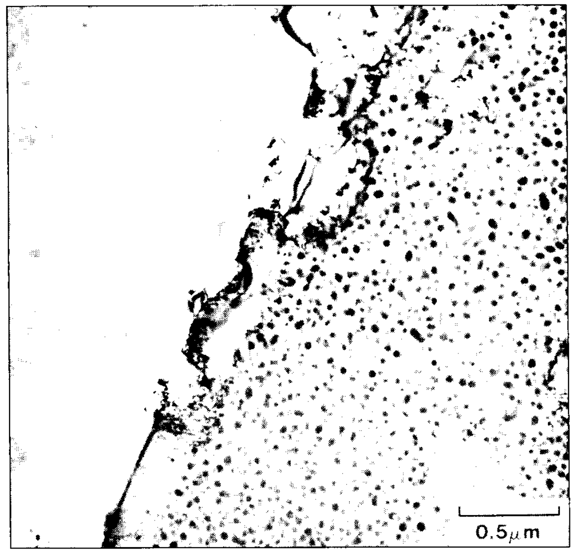

Figure 7. A micrograph of Al-3Li-0.5Mn in peak-aged condition. (from Vasudevan[16]).

```
{{ Al-3Li-0.5Mn-pa
    MEMBER-OF: experimentAl-Li-Mn-series
    MICROSTRUCTURE: Al-3Li-0.5Mn-pa-strc
    ADDITIVES:
        Li
            nominal-percent: 3.0
            unit: weight-percent
        Mn
            nominal-percent: 0.5
            unit: weight-percent
    PROCESS-METHODS:
        cast
            class: direct
        solution-heat-treat
            temperature: 1020
            time: 30
        stretch
            percent-stretch: 2
        age
            time: 48
            temperature: 400
            level: peak
            class: artificial }}
```

Schema 2. A representation of Al-3Li-0.5Mn in peak-aged condition.

Schema 2 includes the following information:

- The base element is Al, by inheritance from the

```
{{ Al-3Li-0.5Mn-pa-strc
     MICROSTRUCTURE-FOR: Al-3Li-0.5Mn-pa
     STRUCTURE-ELEMENTS:
          grain
               size:
                    length: 415
               aspect-ratio: 4
                    alignment: rolling-direction
               texture: copper
                              volume-fraction: 0.02
                         brass
                              volume-fraction: 0.02
                         S
                              volume-fraction: 0.02
                         cube
                              volume-fraction: 0.70
                         Goss
                              volume-fraction: 0.24
               recrystallization-level: 100
               phase: alpha-Al-Li
               structure-elements: precipitate
                                        phase: Al3-Li
                                        size: 0.03
                                                  probability-distribution: log-normal
                                                  aspect-ratio: 1
                                        distribution: uniform
                                        volume-fraction: 0.23
                                                            local-volume-fraction-distribution: log-normal
                                        misfit-strain: 0
                                       dispersoid
                                        phase: Al6-Mn
                                        size: 0.2
                                        aspect-ratio: 3
                                        geometry: rod
                                                  length: 0.3
                                        volume-fraction: 0.005
                                        misfit-strain: high
                                       dislocation
                                        type: mixed
                                        element-density: low
          grain-boundary
               phase: alpha-Al-Li
               angle: high
               impurity: Na K
               PFZ-zone: 0.25
               structure-element: dislocation
                                        type: mixed
                                        element-density: high
                                       precipitate
                                        phase: AlLi
                                        aspect-ratio: 1
                                        geometry: spheroid
                                                  diameter: 1
                                        volume-fraction: 0.04
```

Schema 3. The microstructure of Al-3Li-0.5Mn in peak-aged condition.

experimentAl-Li-Mn-series, which is an instance of the alloy family.

- The microstructure is described in schema Al-3Li-0.5Mn-pa-strc.
- The alloying elements are Li and Mn as specified in the ADDITIVES slot.
- The process methods are, in order, (1) cast, (2) solution-heat-treat at 1020 degrees Fahrenheit for 30 hours, (3) stretch two percent, and (4) age for 48 hours at 400 degrees Fahrenheit, which achieves peak aging.

Vasudevan et al. describe the microstructure — referred to in their paper as "Figure 1(b)" — as follows: "Figure 1(b) shows the microstructure in the peak-aged alloy (condition B), where the strengthening matrix δ' precipitates are seen together with coarse-grained boundary δ precipitates; these are seen as white regions surrounded by dislocations . . . and a δ' precipitate-free zone (PFZ) 0.5μm wide, which has given its solute up to the grain boundary δ."[16]

Schema 3 shows the microstructure representation for the alloy used in ALADIN. Most of the lithium is in the form of either δ' (Al3-Li) precipitates inside grains, or δ (AlLi) particles on the grain boundary. In addition, the grain boundary has a PFZ, with $MnAl_6$ dispersoids distributed throughout the grains. Treating grain interior and grain boundary as separate microstructural elements permitted the association of δ particles and PFZ with the grain boundary, a crucial feature in this microstructure. Characteristics of the microstructure (recrystallization, high-angle grain boundaries, elongated grains parallel to the rolling direction, and low dislocation density) are also represented.

The schema representation — not limited to characteristics apparent on a micrograph — includes quantitative information. The representation's recursive nature (each microstructural element could contain any other microstructural element, even one of the same class) enables metallurgists to represent a broad-range microstructure — an important point. For example, suppose the solution-heat-treated alloy has subgrains inside each grain, and that each subgrain consists of several cells separated by dislocations. A metallurgist could represent such a structure as grains with high-angle boundaries that contain small grains with low-angle boundaries; these small grains contain even smaller grains (or cells) with low- or medium-dislocation density of boundaries. Because grains at each "level" can have various microstructural elements, a metallurgist can represent microstructures with an adequate degree of detail.

Many microstructural elements are associated with a phase (a physically homogeneous part of a material system, characterized by its composition and crystal structure). A phase diagram is a graph of phase stability areas, with composition and environment variables (for example, temperature or pressure) as coordinates.[17] In alloy design, equilibrium phase diagrams define phase regions as a function of temperature and composition, and are referenced when selecting heat treatment temperatures. Occasionally, nonequilibrium phase diagrams are also available that define aging temperatures required to produce metastable precipitates in the material.

Region boundaries are sometimes defined by known thermodynamic equations, which ALADIN can store and evaluate. More often, however, boundaries are determined experimentally. In such cases, each region of an n-dimensional phase diagram can be defined in ALADIN as the union of $(n+1)$-point lattices in n-dimensional space.[5]

Models. Theoretically, we should be able to determine an alloy's properties from its microstructure alone. Practically, the theories are incomplete — requiring the addition of empirical knowledge to fill the gaps. We can view a model as a function that maps from a domain to a range; for example, from microstructure to properties. Due to incompleteness, a model is actually a set of partial functions defined across subsets of the domain. Furthermore, due to uncertainty in these models, domains of partial functions may overlap and map onto different ranges. Consequently, we must represent not only the partial function, but its domain and the credibility of its result. A set of partial functions and associated information (regarding how to choose a function to apply at each step in the design process) can be called a knowledge source. As such, they resemble knowledge sources in Hearsay-II,[18] where stimulus-response frames define a knowledge source's invoking pattern and possible contribution to the interpretation task.

The knowledge base evaluation. The representation defined below achieves most of the goals established at the beginning of this project. This representation is certainly adequate for describing relevant metallurgical information about aluminum alloys and for making necessary distinctions. Indeed, some evidence indicates that the representation can handle non-aluminum materials. We asked two metallurgists — Alcoa's J.L. Murray, and M.A. Przystupa (who is now at UCLA) — to select and describe the microstructures of non-aluminum-based materials, including martensitic steel. We then successfully represented all essential microstructural features of these materials in the ALADIN system. This test revealed no limitations in the representation system. The ALADIN representation provides flexibility with respect to the amount and level of detail, especially in microstructure descriptions. This flexibility is important, since microstructure evaluation techniques vary and details are often not available.

However, such expressive power has a cost. Searching through the nested levels of metaschemata associated with a detailed microstructure description can be computationally expensive. This does not present a problem in our domain. But if this knowledge base were applied to time-critical problems, we would need to evaluate the trade-offs between expressive power and performance. Some materials scientists believe that future products will be specified by microstructure, and that manufacturing facilities will be required to monitor and control those features. In such scenarios, a computer-based representation would be desirable to support real-time control.

The amount of labor required to interpret micrographs and to store essential features in the ALADIN system is, perhaps, a more serious drawback. These tasks, requiring the expertise of someone with an advanced degree in materials science, can be tedious and time consuming. Developing recognition software to automatically interpret micrographs would facilitate the development of a more complete knowledge base in ALADIN.

The problem-solving architecture

Alloy design problems begin by specifying desired physical and mechanical properties of the material to be created, expressed as constraints on these properties. The designer's objective is to identify chemical elements that can be added to pure metal, showing appropriate amounts as a percentage, and processing methods that can be employed to yield an alloy with the desired characteristics. Designers normally use a line of reasoning that resembles goal reduction through a hypothesize-and-test cycle — starting with abstract choices on microstructure, composition, and processing — and proceeding towards a final determination of percentages for each additive and specification of processing practices to be employed.

Microstructure models provide a powerful guide for the design process, since they constrain composition and processing decisions. If metastable precipitates must be present, for example, then (1) the percentage of additives must be constrained below the solubility limit, (2) certain heat treatment processes must be applied, and (3) aging times and temperatures must be constrained within certain numerical ranges.[19]

Similarly, concepts such as solid-solution hardening and interface-boundary strengthening are abstractions referring to mechanisms that can involve a range of additives and process methods, but also lead to narrower subsequent qualitative and quantitative choices. Sooner or later, depending on the design task's nature and the designer's style, a known material is selected as a base line or starting point. The designer then alters the known material's properties by making changes to composition and processing methods, estimates effects of these changes on various physical properties, and identifies discrepancies to be corrected in a later iteration.

However, we must consider several complexities. First, due to the existence of multiple (but incomplete) models, knowledge is often applied opportunistically. For example, it is rarely feasible to predict properties quantitatively from microstructure. But given the class of microstructure, we can often use semi-empirical models to predict properties from composition and processing parameters.

Second, strategies vary among individuals; for example, when selecting the baseline alloy to begin their search, some designers like to work with commercial alloys. Others prefer experimental alloys produced in a highly controlled environment. Still others like to begin with a commercially pure material, and design from basic principles.

Third, when searching for alternatives to meet target properties, some designers construct a complete model of the microstructure (a model that will meet all properties); then, they identify composition and processing options. Other designers 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 avoid microstructure reasoning altogether by using direct relationships between decision variables and design targets.

Fourth, all designers occasionally check their partial plans by estimating primary and secondary effects of fabrication decisions on structure and properties. The frequency of this activity, and the level of sophistication among estimation models, varies among designers.

We specified the following requirements for ALADIN's architecture: The system needs (1) facilities that integrate multiple incomplete design models, (2) a range of control strategies that incorporate individual user preferences, and (3) mechanisms that control search space size and complexity.

Planning and the design process. We designed the ALADIN architecture to support opportunistic reasoning (at different levels of abstraction) across multiple design spaces, and chose a multispatial reasoning architecture akin to a blackboard model.[18,20]

The architecture has five spaces:

(1) Property space — The multidimensional space of all alloy properties (for example, tensile strength, ductility, and fracture toughness);

(2) Structure space — The space of all alloy microstructures;

(3) Composition space — The space in which each dimension represents a different alloying element (Cu and Mg, for example);

(4) Process space — The space of all thermomechanical alloy-manufacturing processes; and

(5) Metaspace — The planning space that directs all processing. The metaspace stores knowledge about the design process and control strategies. Planning takes place in this space, in that goals and goal trees are built there for subsequent execution.

These spaces are subdivided into levels corresponding to different degrees of detail — from abstract qualitative concepts to numerical quantities. Activity is generated on different planes and levels, resembling Stefik's Molgen system.[21] Planes contain one or more spaces.

ALADIN's planes are

(1) The meta (or strategic) plane, which plans the design process and establishes sequencing, priorities, and the like;

(2) The structure-planning plane, which formulates targets at the phase and microstructure level to realize desired macroproperties; and

(3) The implementation plane, which encompasses chemical composition, plus thermal and mechanical processing.

Search activity seeks to reduce the number of outstanding goals through a hypothesize-and-test paradigm. Partial models propose and verify hypotheses at all abstraction levels connecting two or more knowledge spaces.

Continued on page 65

For example, models linking structure and composition can propose alloying elements in the composition space that yield a structure specified in the structure space. In the opposite direction, models can predict properties of a proposed composition by checking whether a phase change would occur in the structure space.

Qualitative and quantitative levels of the structure, composition, and processing spaces are activated (as appropriate) to generate hypotheses specifying design variables. Hypotheses generated on other planes and levels constrain and guide the search for new hypotheses. An existing qualitative hypothesis suggests the generation of a quantitative hypothesis. Compositional additives can produce certain microstructure elements. Specific processes with the composition, restricting available choices, can produce others.

Ideally, alloy design begins in the structure space, where decisions are made on microstructural features that imply desirable properties. Thereafter, these decisions are implemented in composition and process space. Since models are incomplete, the appropriate execution sequence is not known beforehand. Consequently, from the opportunities available, the system must select the most appropriate model at each decision point.

The metalevel planner determines the degree of opportunism exhibited by the system. Its basic cycle

(1) Selects a hypothesis to extend (based upon its credibility),
(2) Evaluates that hypothesis relative to target properties,
(3) Chooses a property to improve,
(4) Picks a model that will optimize the chosen property by refining the hypothesis into new hypotheses, and
(5) Determines how well the new hypotheses meet the selected property goal.

In practice, sequencing among these steps is flexible. For example, selection from among a set of new hypotheses often requires that they be evaluated in detail. Decisions about sequencing at this level are made in the metaspace. The hypothesize-and-test cycle contains a reasoning sequence based on causal relations (represented by linking lines in Figure 8). To evaluate a current hypothesis, for instance, the system determines the effects of composition and process decisions on microstructure. These microstructure estimates are then used to determine the alloy's physical properties. When generating a new hypothesis, on the other hand, causal relations are examined in the reverse direction. Once it is given the target, the system identifies physical properties, microstructure, and then composition and process alternatives. When microstructure knowledge is not available, the system — utilizing the existence of several models — may search for weaker models that bypass the microstructure plane (process-property relations, for example).

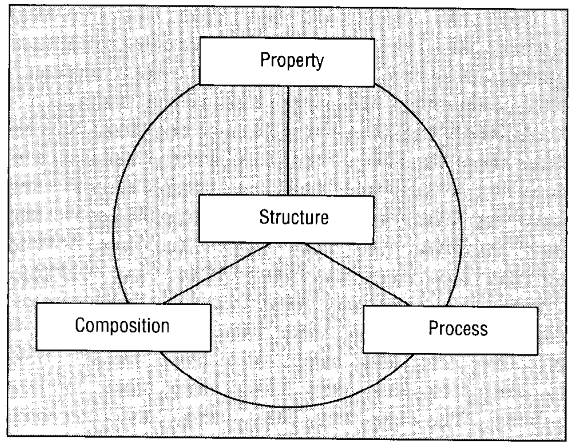

Figure 8. Spaces of domain knowledge.

In its basic hypothesize-and-test cycle, the planner selects goals to satisfy and hypotheses and models with which to satisfy them. It uses several types of information in making such decisions, including

- The search status,
- The solution process history,
- Strategic-alternative constraints, and
- The effectiveness of various strategic alternatives.

The search status is characterized by constraints, hypotheses, and estimates that have been created. These indicate what problems remain. The goals retain solution process history. Given these information sources, constraints on control alternatives are easily represented in rule form. Some examples follow:

If numerical decisions regarding composition
and process have not yet been made
Then quantitative evaluation models cannot
be applied.

If decisions regarding what processing steps
to use have not yet been made,
Then it makes no sense to reason about
temperatures and rates.

Based on heuristic knowledge obtained from metallurgists, the system also has a notion of which strategies will have the greatest impact on search. For instance,

If it is possible to reason about microstructure,
composition, or process,
Then microstructure reasoning is preferred.

If many fabrication alternatives to meet a single
target have been identified,
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 inadequate and the system employs a least-commitment strategy minimizing goal interaction.[22]

ALADIN begins in the metaspace and frequently returns there for new direction. When the metaspace is activated, strategy rules identify reasonable activities and create top-level goals with appropriate context, space, and level information for those chosen activities. Several alternative strategies are often possible at any point during search, and users are offered a menu of possibilities. The system recommends the strategy it estimates most effective. After users accept or override this selection, metarules expand goals by creating more detailed subgoals. These goal trees constitute a near-term, partial plan for accomplishing the design task.

When the metaspace can no longer define goals with current information, control is transferred to the domain space, which processes goals until their success or failure is determined. When no further progress can be made in domain spaces, control returns again to the metaspace, where a new series of goals is defined. These goals are often associated with improving and refining design choices already made. Alternation between metaspaces and domain spaces continues until the problem-solving process is complete.

Within the metaspace, we integrated several design strategies (obtained from different people) into a single system. As a result, ALADIN can develop many solutions to a single problem by applying different approaches. Flexible user control enables metallurgists to experiment with different strategies. For example, users can explore solutions arising from the application of hybrid strategies not usually applied to single problems.

Reducing search complexity. The partial models available to ALADIN can generate a large set of alternative designs. The search space is infinite, in fact, since some design parameters vary along a continuous spectrum. The system uses several domain features to reduce search complexity. The space is somewhat sparse in feasible solutions, and strong heuristics are available to direct search in promising directions. For example, of 103 elements in the periodic chart, only about two dozen are sufficiently soluble in aluminum to be considered as alloying additions. Of those two dozen, we can use only seven as major alloying elements.

Furthermore, to achieve ALADIN's goal of assisting alloy designers and speeding up alloy design, it is not necessary that the program generate feasible or optimal choices in all cases. To generate a family of alternatives that is likely to come close to meeting design goals, when tested, is sufficient.

To reduce search complexity further, ALADIN employs several techniques that are described in AI literature, including

- Hierarchical search,
- Least-commitment search, and
- Constraint-directed search.

Hierarchical search. Many spaces are separated into levels (higher levels being abstractions of lower levels). In the composition space, for example, ALADIN uses the highest level to specify whether an element is to be added or not. Lower levels specify the amounts of an element to add. In the structure space, the highest level specifies the phase type — while the next level down specifies types of microstructural elements present in the microstructure. Still lower levels may contain quantitative information about the microstructure. Planning begins with decisions made at the abstract level — decisions that gradually become more precise, enabling global consequences of abstract decisions to be evaluated before effort is spent in detailed calculations.

Metallurgical models act as partial filters on design decisions. That is, when decisions regarding composition and process are made, they are filtered through the structure space to predict their effects in the property space. The filter is partial, since an appropriate metallurgical model may not exist in all cases. Consequently, microstructure decisions form an abstract plan that cuts down the number of alternatives in composition and process spaces. In this way, the role of the microstructure differs in some respects from abstract planning, as Sacerdoti describes.[23] The main differences are that

- Microstructure concepts are distinct from composition and process concepts, and not merely less-detailed descriptions;
- The microstructure plan is not part of the final design, since an alloy can be manufactured with composition and process information only; and
- The microstructure domain is predefined by metallurgical expertise, and not defined during ALADIN's implementation or execution.

These differences lead to the following contrasts with a Molgen-like system:

- Instead of one plan hierarchy, ALADIN has three — structure, composition, and process — each of which contains abstraction levels.
- Since structure decisions do not necessarily have the highest "criticality" (as defined by Sacerdoti), opportunistic search is important.
- The effect of abstract hypotheses is more complex, because decisions in the structure space cut search by constraining the choice of both composition and process hypotheses. The existence of more than one level in each space also introduces new interaction types.

Least commitment. Since most design decisions affect more than one target property, it is often inappropriate to

```
{{ AGE-TEMPERATURE-DEFAULT-MODEL
        IS-A: model
        MODEL-OF: temperature
        CREDIBILITY: 0.2
        DOMAIN: (type class artificial)
        TEMPERATURE: 400 }}
```

Schema 4. A schema representation of a typical-aging-temperature model.

```
{{ AGE-TEMPERATURE-MODEL
        IS-A: model
        MODEL-OF: temperature
        CREDIBILITY: 0.9
        DOMAIN:
            range: (or (type class natural) (type class artificial))
            cardinality: (1 2)
        TEMPERATURE:
```

Schema 5. A general age temperature model.

make precise commitments early in the design process. Furthermore, the complexity of design decision interdependence makes simple concepts of goal protection inefficient.[8] Hence, ALADIN employs a least-commitment strategy that minimizes goal interaction. The system expresses quantitative hypotheses as constraining inequalities, which are kept as broad as possible. Hypotheses are refined by posting additional constraints that reduce the region for design variables, thereby reducing the need for backtracking when selecting values.

ALADIN's domain lends itself readily to this technique: The system can use qualitative reasoning to determine what constraint should be refined (and most numerical variables admit to value ranges).

Constraints. Although it focuses on a single property in the hypothesize-and-test cycle, ALADIN uses constraints to relate properties. As a result, the system need not spend time exploring decision paths that would adversely affect other properties; constraints are represented as Lisp expressions involving any function or variable evaluating to a non-negative result. Formulas for density and modulus immediately yield constraining equations. ALADIN can obtain constraining equations for other properties by regression in the alloy database. Some variables are not easily quantifiable, but have an indirect impact on generated constraints. For example, temper information is used to select alloys during regression. Phase diagrams — heuristic rules involving phase boundaries and solubility limits — are another source of constraining equations. During evaluation, the system tests all constraints to verify that the design is acceptable — not only with the current goal, but with the other property goals.

In some cases, constraints are also used to generate decisions. In such instances, a variant of Hadley's gradient method is used to find a feasible point for a system of nonlinear inequalities.[24]

Model-based inference

Experience accumulated during our interaction with metallurgists, and insights gained during our work with ALADIN, suggest a need for an alternative operational mode. Typical users of an alloy design system will, for the foreseeable future, be metallurgists with considerable expertise in at least some aspects of alloy design. Each metallurgist has an individual style (and, often, firm opinions) on what approach should be taken. Some metallurgists prefer a system that leaves the top control to users, but assists design by suggesting alternatives. That is the purpose of the design assistant mode, in which metallurgists guide search in the direction they want. Hypothesis elaboration is also put under user control by providing a set of models, with which users can derive new information.

The design assistant mode applies the general and powerful notion of models. We created a schema-based model representation and a domain-independent inference engine that invokes models to infer values of attributes in schemata.[25] ALADIN uniformly represents domain-dependent information, facts, qualitative and quantitative models — as well as much domain-independent control knowledge — in schema form.

Reasoning involves inferring attribute values in existing or newly created schemata. If we can obtain acceptable values through simple retrieval, with or without inheritance, the values are considered known and no model need be invoked. Otherwise, a value will be inferred (if possible) through a search for the "best" model that yields an acceptable result. Models are selected in three stages. First, the domain of model validity is determined. The validity domain is a subset of all schemata, specified with Knowledge Craft's restriction grammar.[12] For example, Schema 4's DOMAIN attribute limits the use of that model to temperature schemata of "type class artificial." Second, valid models are ranked according to their credibility. Third, the value generated by the model has to satisfy range and cardinality restrictions; for example, Schema 5's DOMAIN must be one or two schemata of "type class natural" or "type class artificial."

Searching and ranking models — as well as determining domain, credibility, and range — are inferences that the control models can perform. ALADIN has a set of domain-independent control models that can be augmented and superseded by domain-dependent control models whenever appropriate.

The simplest use of a model, to infer a specific slot value in a schema, is to take the value from the same slot

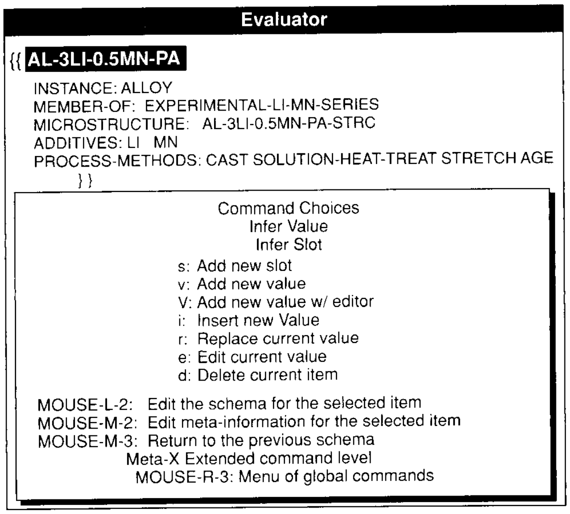

Figure 9. The Infer Value and Infer Slot items can be selected.

of a similar schema (which then becomes a model or analog of the first schema). For example, if we want to determine the aging temperature for Schema 2's alloy, we can use knowledge about the typical temperature for artificial aging as a model (see Schema 4) and assume that the temperature is 400 degrees Fahrenheit. The Schema AGE-TEMPERATURE-DEFAULT-MODEL is declared to be a temperature model through the MODEL-OF relation.

We can introduce algorithmic and numerical models through procedural attachment; for example, we can generate a value by attaching a piece of code to an attribute. Schema 5 is a model that invokes a procedure specified by the AGE-TEMPERATURE-MODEL-PROCEDURE schema.

This model inference system can be viewed as an extension of more conventional schema representation features, and is implemented as a function (infer-value) to be used instead of the function (get-value) provided by Knowledge Craft's schema representation system. Mechanisms for searching and selecting attribute models make it possible to distinguish cases based on complex criteria (numerical relationships, for example). The range and cardinality checks on inferred values implement a simple backtracking feature. Successfully inferred values are optionally stored in the schema, with meta-information on their source. Hence, if the same call to the infer-value function is repeated, the value will be obtained by simple retrieval. This is also true regarding input data

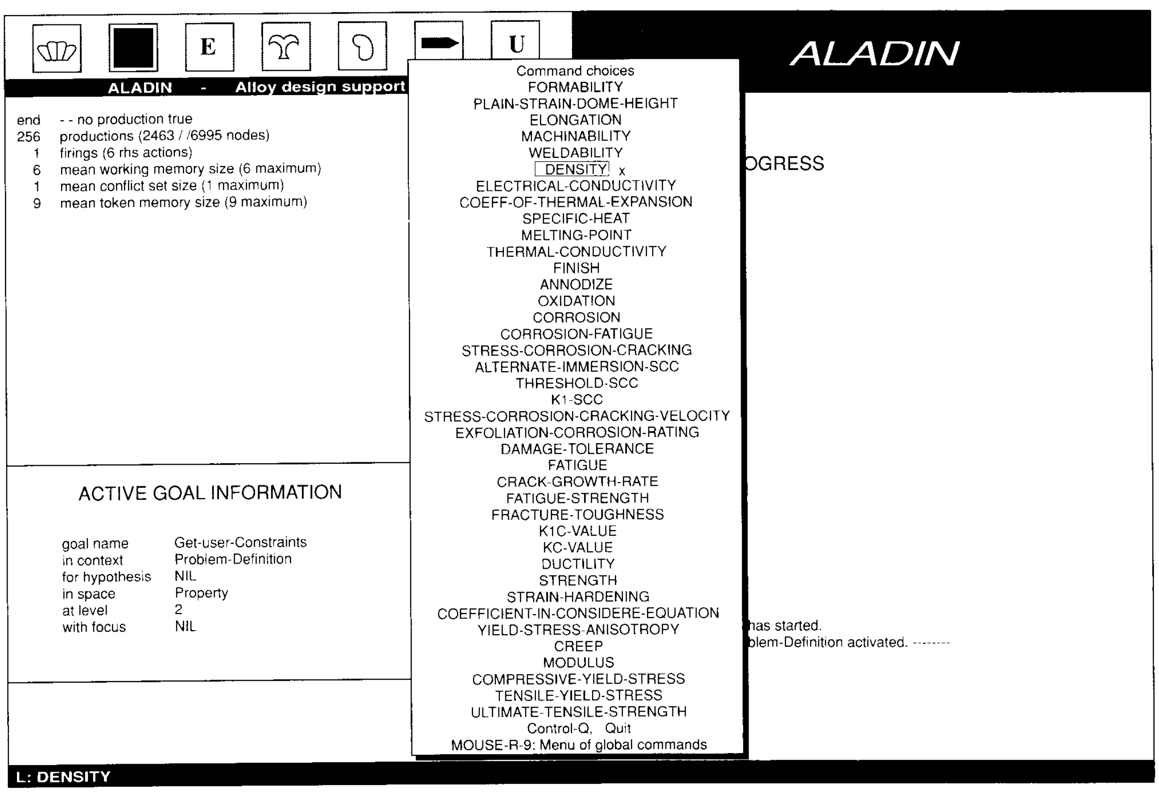

Figure 10. The user is presented with a menu of properties to constrain and will be given the opportunity to limit each selected property quantitatively.

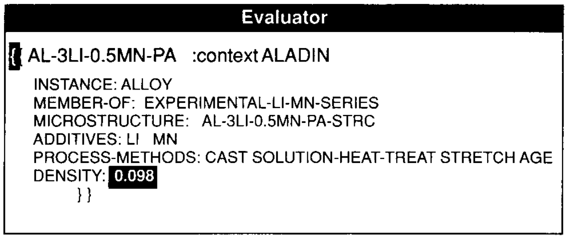

Figure 11. The Infer Value item gives the result.

and intermediate results obtained by recursive calls to the infer-value function, either by the selected model or the infer function itself.

If we adopt a convention to store only facts as regular values, and to represent default values as models, then this architecture provides a clean cut between defining properties and default properties. This addresses a well-known ambiguity in frame-based knowledge representation systems.[26]

Resembling Knowledge Craft's schema editor, the design assistant enables users to invoke models and to enter information (including simple editing commands) in an interactive environment. Figure 9 shows the menu of top-level commands. Selecting the Infer Slot command generates a menu of slots (that is, attributes) that are appropriate in the displayed schema. In this case, the menu would resemble the one in Figure 10.

Selecting an attribute introduces it in the displayed schema. The Infer Value command activates the inference engine described above, and inserts the resulting value. Figure 11 shows the result of inferring the density.

Basic system behavior

ALADIN runs on a Symbolics Lisp machine, under Genera 7.1, within the Knowledge Craft 3.1 environment,[12] at a suitable speed for interaction with expert alloy designers. The design run outlined in this section (Figures 10, 12, and 13 represent an ordered sequence of snapshots from an ALADIN run) takes about half an hour and involves considerable interaction with users, whose choices influence the outcome. The system has reached the advanced-prototype stage and can assist in the design process, especially as a knowledge base and design evaluator (two of the three main functions we set out to develop — the third mode involves independent design and discovery). Presently, however, the knowledge base focuses

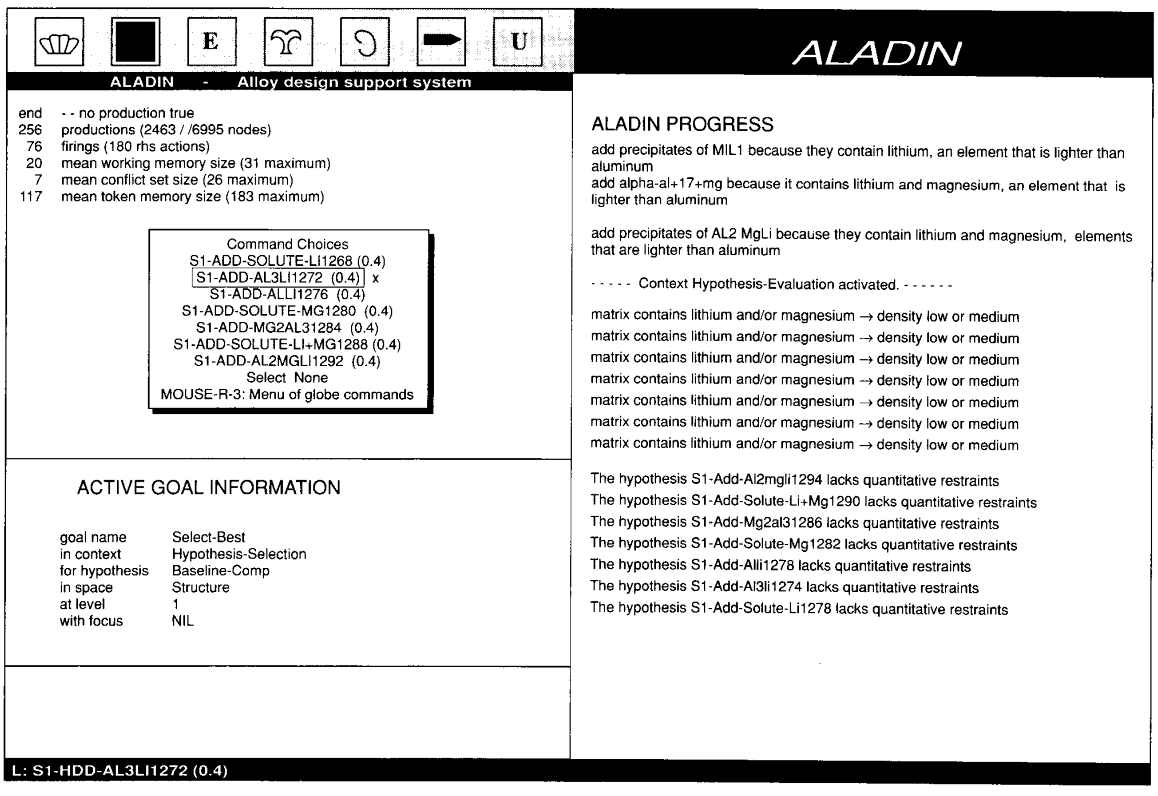

Figure 12. A qualitative evaluation (which notices that the addition of light elements implies low density) is performed. A menu of hypotheses enables users to force a selection.

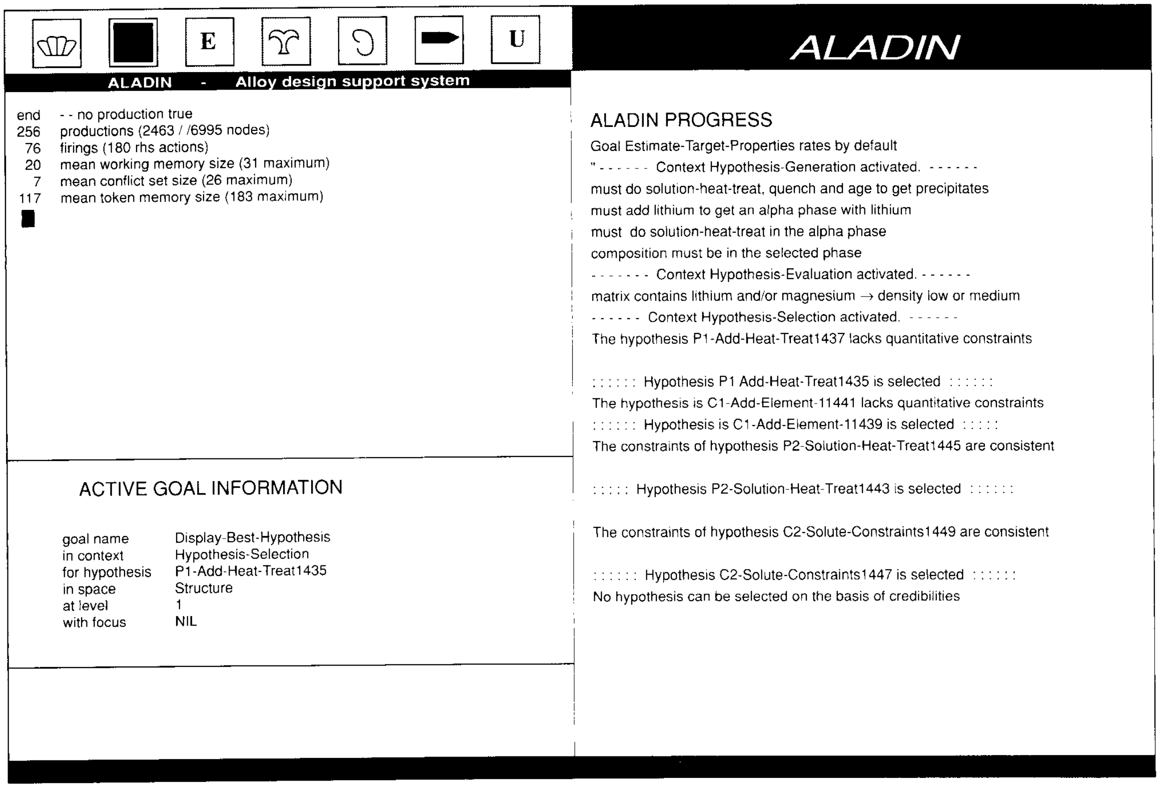

Figure 13. Quantitative constraints are generated (and checked for consistency, if present).

on narrow areas of alloy design and contains expertise on only three additives, two microstructural aspects, and five design properties. Some heuristic rules are ad hoc, rather than being integrated into the strategy-planning-implementation hierarchy. We have dealt in depth only with ternary alloys. But these restrictions were by our own choice, so that we could concentrate on selected areas of greatest importance to our expert informants and sponsors. Within these restrictions lie numerous commercially important alloys, whose rediscovery by ALADIN would be a major milestone.

In its independent design mode, ALADIN uses goal-based reasoning. The system's first goal is to obtain a target for the desired alloy. A design target is generally described in terms of values on various physical properties, although some design problems can be specified in terms of microstructure constraints. Physical property targets can be associated with any of the attributes listed in Figure 10.

During its setup phase, ALADIN also determines the material's application (aerospace, let's say, or packaging) and the product form required (sheet or extruded rod). Using this information, the system prioritizes property goals, defines an appropriate search strategy, and selects a starting point for design search. It makes each of these decisions by applying its knowledge base of metallurgical design techniques and guidelines. ALADIN can also express individual user preferences through its menu-driven interface, and can use these preferences to modify default guidelines.

Once the problem is defined and search is set up, the system enters a cycle of hypothesis generation, selection, and evaluation. Figure 14 shows a generation step in which alternative microstructural phases (which may achieve property targets) are retrieved from the knowledge base. These phases — solid-solution Al-Li, AlLi precipitates, Al_3Li precipitates, solid-solution Al-Mg, Mg_2Al_3 precipitates, solid-solution Al-Li-Mg, and Al_2MgLi precipitates — all contain Li and/or Mg to reduce the starting alloy's density. After each generation step, the system evaluates alternatives relative to targets (as Figure 12 illustrates). Proposed alloy descriptions are incomplete during early design iterations, containing only qualitative microstructural choices. Hence, only approximate evaluations can be made for each choice. Nevertheless, this is sufficient to uncover major differences in options and to point search in a promising direction.

During later design iterations, the system incrementally develops a hypothesis tree containing quantitative refinements of early decisions and constraints on

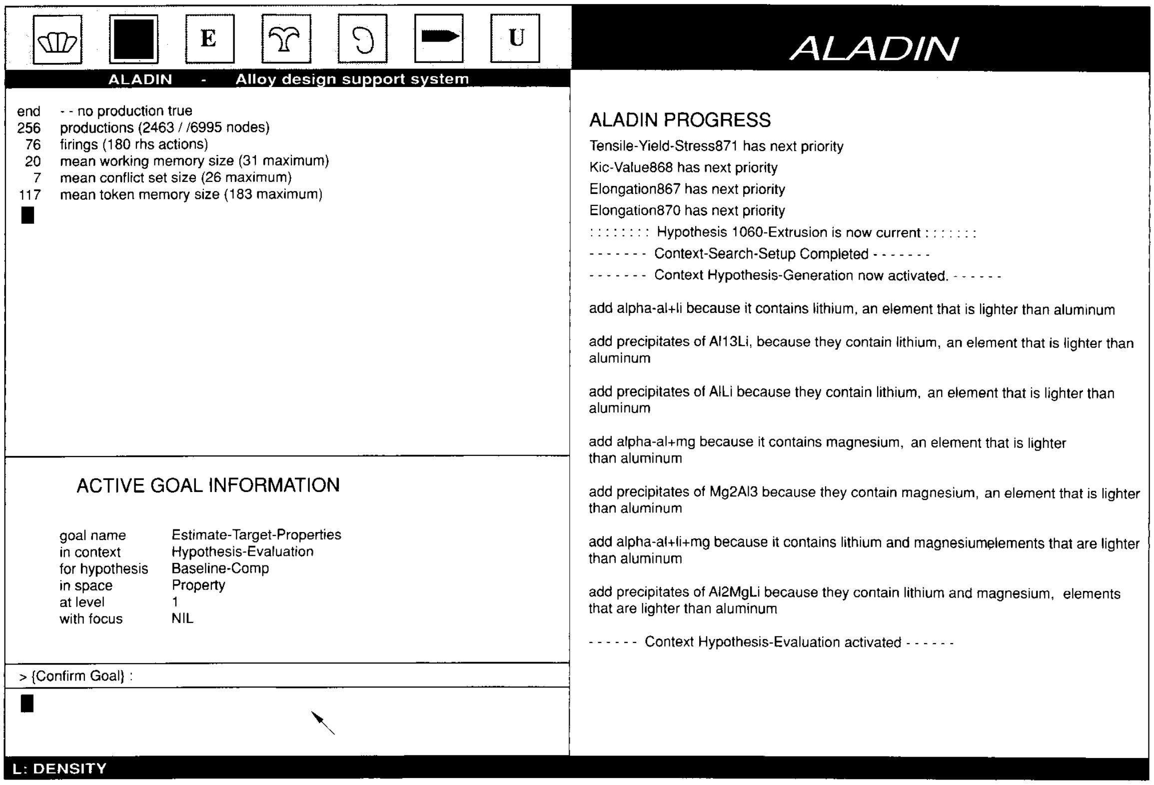

Figure 14. In this example, six hypotheses on the microstructural composition are generated — all involving light elements in response to a low-density target.

composition and processing propagated from microstructural constraints. For example, to produce a phase of Al_3Li precipitates in the starting material, Li should be added at a level above the room temperature solubility limit, and the product should be solution heat treated at a temperature above the phase diagram's solvus line.

Because the search for a new alloy is usually driven by product requirements, designers generally have an application in mind. The system uses this information to select a design strategy. Since ALADIN pursues one target at a time, it must prioritize targets.

The system uses its database of known commercial and experimental alloys for qualitative and quantitative comparisons with the target — comparisons best made between alloys of similar product forms. In Figure 12, only a qualitative evaluation is performed. To form a basis for selection, the subsequent selection phase assigns credibilities to alternative hypotheses. In such cases, no quantitative constraints are available that could affect selection.

The hypothesize-select-evaluate cycle adds design details incrementally and builds a hypotheses tree (as Figures 12 and 13 show). A design session produces a partial description of an alloy, using the knowledge representation of the alloy database.

Alloy design has been thought to require too high a degree of creativity and intuition for automation. However, we have found that hypothesize-and-test cycles, abstract planning, and rule-based heuristic reasoning can reproduce a significant portion of the reasoning used by human designers on prototypical cases. Nevertheless, the knowledge required to solve real-world problems is enormous. Metallurgists have suggested that future efforts focus on more specific aspects of material design. Methods developed here form the necessary basis and framework for such efforts; the current ALADIN system has approximately 2400 schemata, 250 Knowledge Craft OPS rules, and 200 Lisp functions.

This research has

- Provided a representation in which multiple partial models can be represented declaratively;
- Formulated an architecture in which incomplete and even inconsistent models can be integrated in the design process;
- Satisfied multiple interacting goals by determining least-commitment constraints;
- Developed a framework and applied a set of techniques that permit the effective coupling of symbolic (qualitative) and numerical (quantitative) reasoning,

within a structure containing various representations of information;

• Found ways to reason qualitatively with constraints that are expressed quantitatively; and

• Created an interactive environment where experts can share control of the design process with the system.

As currently practiced, alloy design may involve many iterations over several years. ALADIN's overall goal, as an industrial application of AI techniques, has been to make the alloy design process more productive.[27] ALADIN can achieve significant productivity improvements and speed the discovery of better alloys by (1) making the generation of alloying experiments more systematic, (2) aiding the evaluation of proposed experiments, and (3) enabling individual designers to supplement their own specialized expertise with that of the program, which is a pool of expertise from various sources.

The domain complexity has given us the opportunity to extend the frontiers of AI research. Search in the space of abstract models (in our case, microstructure) has potential applications in other design areas — the design of other metallic or nonmetallic materials and, in general, designs dominated by nongeometrical constraints that require a combination of qualitative and quantitative reasoning.

Our representation of strategic knowledge, with flexible user control, provides a powerful means for combining knowledge from multiple experts into a single system. We hope that these ideas will be useful to developers of future expert systems.

Acknowledgments

This research has been supported by the Aluminum Company of America. We thank our expert metallurgist informants from Alcoa — especially Marek Przystupa, Douglas Marinaro, A. Vasudevan, Warren Hunt, James Staley, Philip Bretz, and Ralph Sawtell. We also thank Cheryl Begandy and Walter Cebulak for project support and direction.

References

1. D. Marinaro and J.W. Morris, Jr., "Research towards an Expert System for Materials Design," in *Artificial Intelligence Applications in Materials Science*, R.J. Harrison and L.D. Roth, eds., Metallurgical Society, 420 Commonwealth Dr., Warrendale, Pa., 1986, pp. 49-77.

2. I. Hulthage et al., "The Representation of Metallurgical Knowledge for Alloy Design," *Artificial Intelligence for Engineering Design, Analysis, and Manufacturing (AIEDAM)*, Vol. 1, No. 3, 1988, pp. 159-168.

3. M.L. Farinacci, I. Hulthage, and M.A. Przystupa, "Acquiring and Representing Knowledge about Material Design," in *Knowledge-Based Expert Systems in Engineering: Planning and Design*, D. Sriram and R.A. Adey, eds., Computational Mechanics Publications, Southampton, England, 1987, pp. 99-114.

4. M.D. Rychener et al., "Integrating Multiple Knowledge Sources in ALADIN, an Alloy Design System," *Proc. Fifth Nat'l Conf. Artificial Intelligence*, MIT Press, Cambridge, Mass., 1986, pp. 878-882.

5. I. Hulthage et al., "The Use of Quantitative Databases in ALADIN, an Alloy Design System," in *Coupling Symbolic and Numerical Computing in Expert Systems*, J.S. Kowalik, ed., Elsevier North-Holland, New York, N.Y., 1986.

6. M.L. Farinacci et al., "The Development of ALADIN, an Expert System for Aluminum Alloy Design," *Proc. Third Int'l Conf. Advanced Information Technology*, Elsevier North-Holland, New York, N.Y., 1986.

7. J.R. Dixon, "Artificial Intelligence and Design: A Mechanical Engineering View," *Proc. Fifth Nat'l Conf. Artificial Intelligence*, MIT Press, Cambridge, Mass., 1986, pp. 872-877.

8. A.A.G. Requicha, "Representations for Rigid Solids: Theory, Methods, and Systems," *Computing Surveys*, Vol. 12, No. 4, 1980, pp. 437-464.

9. E. Hornbogen, "On the Microstructure of Alloys," *Acta Metallurgy*, Vol. 32, No. 5, 1984, p. 615.

10. W.A. Boag, Jr. et al., "Cordial — A Knowledge-Based System for the Diagnosis of Stress Corrosion Behavior in High-Strength Aluminum Alloys," in *Artificial Intelligence Applications in Materials Science*, R.J. Harrison and L.D. Roth, eds., Metallurgical Society, 420 Commonwealth Dr., Warrendale, Pa., 1986, pp. 123-146.

11. W.A. Woods, "What's Important About Knowledge Representation?" *Computer*, Oct. 1983, pp. 22-27.

12. *Knowledge Craft, Version 3.1*, Carnegie Group, 5 PPG Place, Pittsburgh, Pa., 1986.

13. M.S. Fox, "On Inheritance in Knowledge Representation," *Proc. Sixth Int'l Joint Conf. Artificial Intelligence*, AAAI, Menlo Park, Calif., 1979, pp. 282-284.

14. E.E. Underwood, *Quantitative Stereology*, Addison-Wesley, Reading, Mass., 1970.

15. I. Hulthage et al., "The Metallurgical Database of ALADIN — An Alloy Design System," in *Artificial Intelligence Applications in Materials Science*, R.J. Harrison and L.D. Roth, eds., Metallurgical Society, 420 Commonwealth Dr., Warrendale, Pa., 1986, pp. 105-122.

16. A.K. Vasudevan et al., "Fracture Behavior in Al-Li Alloys: The Role of Grain Boundary δ," *Materials Science Engineering*, Vol. 72:L25, 1985.

17. L.H. Van Vlack, *Elements of Materials Science and Engineering*, Addison-Wesley, Reading, Mass., 1975.

18. L.D. Erman et al., "The Hearsay-II Speech-Understanding System: Integrating Knowledge to Resolve Uncertainty," *Computing Surveys*, June 1980.

19. D.A. Porter and K.E. Easterling, *Phase Transformations in Metals and Alloys*, Van Nostrand Reinhold, New York, N.Y., 1981.

20. B. Hayes-Roth, "A Blackboard Architecture for Control," *Artificial Intelligence*, Vol. 26, 1985, pp. 251-321.

21. M.J. Stefik, "Planning with Constraints (Molgen: Part 1)"; "Planning and Metaplanning (Molgen: Part 2)," *Artificial Intelligence*, Vol. 16, 1981, pp. 111-170.

22. R. Wilensky, *Planning and Understanding: A Computational Approach to Human Reasoning,* Addison-Wesley, Reading, Mass., 1983.

23. E.D. Sacerdoti, "Planning in a Hierarchy of Abstraction Spaces," *Artificial Intelligence*, Vol. 5, 1974, pp. 115-135.

24. G. Hadley, *Nonlinear and Dynamic Programming,* Addison-Wesley, Reading, Mass., 1964.

25. I. Hulthage, "Reasoning with Models in Artificial Intelligence," in *Methodologies for Intelligent Systems IV*, Z.W. Ras, ed., Elsevier North-Holland, New York, N.Y., 1989.

26. R.J. Brachman, "'I Lied about the Trees' Or, Defaults and Definitions in Knowledge Representation," *AI Magazine*, Vol. 6, No. 3, 1985, pp. 80-93.

27. G. Khermouch, "Alcoa Vigorously Pushes an Array of AI Projects," *American Metal Market*, July 2, 1987.

Ingemar A.E. Hulthage is now a scientist at the Carnegie Mellon Research Institute. From June of 1984 through June of 1990, he was a faculty member of the Carnegie Mellon University Robotics Institute. He earned his PhD in theoretical physics from Stockholm University in 1981. Before joining the Robotics Institute, he did physics research at CMU, SUNY at Stony Brook, Niels Bohr Institute in Copenhagen, and Stockholm University — and has written or co-authored eight publications in theoretical physics. He has been active in several projects on the intelligent processing of materials, is now involved in research on sequencing the forging of steel billets, and has taken a special interest in coupling symbolic and numerical computing in expert systems. He is a member of AAAI and APS.

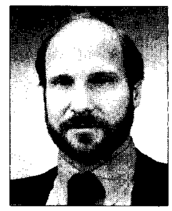

Mark S. Fox received his BSc from the University of Toronto in 1975, and his PhD from Carnegie Mellon University in 1983. He joined CMU's Robotics Institute as a research scientist in 1979, was appointed director of CMU's Intelligent Systems Laboratory in 1980, and cofounded Carnegie Group Incorporated in 1984. CMU appointed him associate professor of computer science and robotics in 1987 and, in 1988, named him director of the Center for Integrated Manufacturing Decision Systems. An *IEEE Expert* Editorial Board member for four years, Fox published an article in *IEEE Expert* earlier this year (February 1990, pp. 8-20). He has published over 50 papers in all, and pioneered the application of AI to factory planning and scheduling problems, project management, and material design.

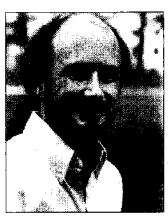

Michael D. Rychener is a member of technical staff at Bell Communications Research, working on automated knowledge acquisition and knowledge management for expert systems. When this article was being done, he was a senior research scientist at CMU's Engineering Design Research Center, applying knowledge-based expert system techniques to engineering design problems. His recent book, *Expert Systems for Engineering Design* (Academic Press, 1988) describes some of these projects. He earned his BA from Oberlin College in 1969, his MS in computer science from Stanford University in 1971, and his PhD in computer science from Carnegie Mellon University in 1976. His doctoral dissertation was one of the first to explore the possibilities of rule-based systems for AI applications. He is a member of the IEEE Computer Society, ACM, and AAAI.

Martha L. Farinacci received her BS in mathematics from Rensselaer Polytechnic Institute, and her MS in applied mathematics from Carnegie Mellon University. Her technical interests include knowledge-based systems, operations research, software design, and numerical analysis. She has spent 12 years in industry, developing applications of mathematics, operations research, and computer science to problems of manufacturing planning, process control, and engineering analysis. At CMU's Robotics Institute and the Alcoa Research Laboratory, she developed AI-based systems to help scientists and engineers with product design problems. Currently a member of technical staff at the Mitre Artificial Intelligence Center in Washington, DC, she designs and evaluates knowledge-based systems for various government agencies.

The authors can be reached in care of Ingemar Hulthage, Carnegie Mellon Research Institute, 4400 Fifth Ave., Pittsburgh, PA 15213.