

# UNDERSTANDING AND AUTOMATING ALGORITHM DESIGN

Elaine Kanf

Schlumberger-Doll Research  
Old Quarry Road  
Ridgefield, Connecticut 06877-4108

## Abstract

Algorithm design is a challenging intellectual activity that provides a rich source of observation and a test domain for a theory of problem-solving behavior. This paper describes a model of the algorithm design process based on observations of human design. The adaptation of that model to automation in the DESIGNER system helps us understand human design better, and the automation process helps validate the model. Issues discussed include the problem spaces used for design, the loci of knowledge and problem-solving power, and the relationship to other methods of algorithm design and to automatic programming as a whole

## 1. The Algorithm Design Task

### A. Design as an Intellectual Activity

Algorithm design is the process of coming up with a sketch, in a very high level language, of a computationally feasible technique for achieving a specified behavior. The design process combines cleverness in problem solving, knowledge of specific algorithm design principles, and knowledge of the subject matter of the algorithm (e.g. geometry, graph theory, physics). When people design algorithms, their design repertoire includes discovery and visual reasoning in addition to the (ideally) disciplined application of problem-solving techniques

Human design is a rich source of ideas for a model of algorithm design. Observing that design process and attempting capture the basic ideas in an automated system both helps us understand how people structure and use their knowledge about design and also validates our observations and model. The DESIGNER project included such a study of human design and an initial version of an automated system [15, 16, 17, 26].<sup>•</sup> The goal of the project is to create an automatic design system that can apply existing design principles as well as exhibit some creativity. The observations of human design are to be incorporated, but the automatic system should take the strengths and weaknesses of both computers and people into account. We are not trying to model human problem-solving behavior as an end in itself.

<sup>•</sup>This paper describes research performed while the author was at Carnegie-Mellon University. The research was supported in part by DARPA and in part by NSF.

<sup>•</sup>The research described here is joint work with Allen Newel and David Steier. Many of the ideas I draw on are theirs, but the opinions expressed here are my own.

The next section of this paper (1.B) presents a sample algorithm design problem (finding the convex hull). Section II summarizes our observations of human designers working on that and on other algorithms. Section III then discusses where the problem-solving power in our model lies: in the ability to search in multiple spaces (relying on knowledge from the domain as well as knowledge about algorithm design), on efficiency knowledge, and on the ability to execute partial algorithm descriptions on examples. Finally, this model of design is compared in Section IV with other approaches to automating algorithm design and to automatic programming as a whole.

### B. A Challenge: The Convex Hull Problem

Consider the problem of finding a convex hull, which has applications, for example, in algorithms for vision and graphics. The problem is this: Given a set of points in two dimensions find a polygon whose vertices are a subset of those points that encloses all the other given points.

Now if I drew some points on a blackboard or piece of paper, you would probably have no trouble sketching their convex hull. (If you need a picture to help you understand the problem, see Figure 1.) Suppose instead I asked you, or an automatic design system, to create an algorithm suitable for (later) encoding as a computer program in a conventional high-level language. Think of sketching out an algorithm in the terms you would use for describing it to a colleague or to a programmer, without worrying about the low-level implementation details. As you work on this problem, observe your problem-solving behavior. Do you write down any formal problem descriptions? In what language? Do you draw pictures? Create a variety of examples or counter-examples? Draw analogies to other algorithms? Draw on general knowledge of algorithm design principles?

How did you convince yourself that your algorithm was correct? Did you design it by applying correctness-preserving transformations? Did you use geometric or other mathematical theorems? Find proofs for conjectures? Test your algorithm on sample data? Explain the algorithm to yourself or a friend in words? Write pseudo-code?

How do you decide when your algorithm is complete? How do you decide when it is good enough? What does it mean to be a good algorithm? Do you know what the run-time or space performance of your algorithm will be? Did you worry about what the distribution of data would be in creating the algorithm? In determining performance?

Now let's find out how some other people solved this problem and see how we might design an artificial intelligence program that could perform the same feat.

## II. Methods for Designing Algorithms

Since the design of complex algorithms is currently best accomplished by human beings, observing their performance would appear to be a profitable starting point for automating the design process. However, since the talents of computers are not those of people, it is reasonable to search for a different method if the goal is total automation of design or a novel mixture of human and machine design. This issue is discussed in Section IV

The model of design presented here is based on the analysis of a set of protocols from approximately fifteen sessions with computer science faculty, graduate students, and undergraduates (A methodology for protocol analysis is described elsewhere [8, 20]) Our designers were independently given the task of creating algorithms to find convex hulls, closest pairs (if point), and intersecting line segments. Several protocols have been analyzed in great detail while the others have been gone over more lightly and used primarily as confirming evidence.

Before summarizing the features of human design, some caveats on the general applicability of the observations are in order (1) We observed the design of individual algorithms whose complexity is due to a requirement for cleverness rather than to the information processing overload of combining an overwhelming number of small but straightforward parts (2) The algorithms depend on applying an appropriate set of operations rather than on designing a specialized data structure (3) Our study did not include any interaction between people and design aids other than pencil and paper or blackboard (However, no one volunteered any feelings that a calculator, computer, or any other automated device would have been of any help in designing their algorithm.) (4) The design sessions we observed were on the scale of hours rather than the months spent by research algorithm designers. Other processes than those we observed may take place in such long time periods

Our observations may be at least partially valid in a wider context despite the caveats. Other researchers have studied the design process in software engineering and have made observations similar to ours [1,12]. Also, there is anecdotal evidence that similar problem-solving techniques are used in the design of algorithms that are highly dependent on clever representations

The processes that we observed our designers draw upon include:

1. *Understand* the problem.
2. *Select* a problem to work on
3. *Plan* a solution around a kernel idea and *refine* or elaborate the kernel structure
4. *Execute* the partially specified algorithm.
5. *Notice* and formulate any difficulties or opportunities.
6. *Verify* that the structure is a solution (i.e. meets its specifications)
7. *Evaluate* the solution (e.g. for efficiency)

After the processes in this collection are summarized the issue of control — how the processes are ordered and evaluations within each step — will be discussed. The explanations draw on all of our observations and those of our colleagues who have studied the design process in software engineering. However, I will give illustrations primarily from the stories of two particular designers from our study D1 and D2, who tried to solve the convex hull problem. Each part of a story is prefaced by the designers name

and number for future reference For example the first step of Designer 1's story is labeled [D1.1]

### A. Understand the Problem

In classical discussions of problem solving [22] one important problem-solving process is understanding the problem, perhaps by listing properties of the objects in question, and considering reformulations of the problem. Some of our designers (but not D1 or D2) did draw a picture of a convex hull (or whatever) early on, which may have led to some un verbalized observation of or reasoning about properties of convex hulls and seemed to have convinced them that they understood the problem.

[D6.11] D6 drew the picture shown in Figure 1.

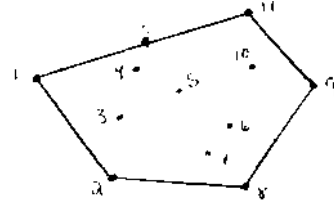


Figure 1 Initial example drawn by {

[D3.1] D3 wondered whether using polar coordinates might not be a useful way to think about the problem.

### B. Plan and Refine Solution

Assuming that a problem specification has been understood, design begins with a *kernel idea* or solution plan, quickly selected from those known to the designer Depending on the designer's background, the idea may vary in sophistication from generate and test to input process output to more complex strategies such as divide and conquer or dynamic programming. The designer lays out the basic steps of the chosen idea and follows through with it unless the approach proves completely inapplicable

[D1.1] D1 had the initial idea that the algorithm should be one that generated all points in the input in some arbitrary order and tested each to determine whether it was on the hull. This had the potential of running in linear time (proportional to the number of input points)

[D2.1] D2 decided to try a divide and conquer algorithm (the special form of divide and conquer in which the input, are divided into subsets, the algorithm is recursively applied to each, and the results are back together)

Which kernel idea do you think will lead to a better algorithm?

After formulating a plan, the designer *refines* the basic steps of the kernel idea. By and large, this elaboration proceeds by stepwise refinement. The designer may lay down the major components, effectively decomposing the problem into subparts, or may add new inputs or assertions about details of the structure. The refinement steps (1) may be suggested by *knowledge* appropriate to the problem and task domain or (2) may be a natural result of attempting to *execute* an algorithm

[D2.2] An example of the application of appropriate knowledge about algorithm design principles is D2's expansion of the notion of using a divide and conquer algorithm into the sequence of steps: divide input point sets into subproblems, find the convex hulls recursively, merge subsolutions back into a convex hull

[D2.3] Furthermore, D2 recognized from previous experience with geometric algorithms that a likely possibility for the divide step of the divide and conquer algorithm was to sort the points by one of the coordinates and find the median

In the absence of the knowledge that suggests the proper refinements, the designers *search* by trail and error, they hypothesize algorithm steps and try them out by executing the partially specified algorithm.

[D1.2] D1 had no idea how to test whether a point was on the hull and decided to try out the algorithm on a specific figure to find the test.

[D1.3] D1 then drew the picture shown in Figure 2.




Figure 2. Initial example drawn by D1.

However, the refinement process is hardly one of pure top-down design.

[D1.4] Such a point-on-hull test didn't reveal itself, but another related test did, and D1 proceeded to modify the hypothesized algorithm to exploit the new test.

Most design falls in between having the correct knowledge and searching. At some steps the designer knows what to do and knows what the implications of the refinement step are; other times, search is required.

[D2.4] D2 did not find the merge step as obvious as the divide step. Do you?

### C. Execution of Algorithms

*Trial execution* of algorithms is often used as a technique for making inferences about the algorithm developed so far. We observed two kinds of execution of partially specified algorithms — one on concrete data (which we call *test-case execution*) and the other on symbolic exemplars (which we call *symbolic execution*). Both forms of trial execution help elaborate the algorithm description by exposing difficulties and opportunities. We found it useful to view execution as a technique for selectively *propagating* constraints (which we call assertions,) by moving them around in the order in which steps of the algorithm are executed. This limits the reasoning that might otherwise be necessary to find contradictions and make inferences

### D. Noticing Difficulties and Opportunities

Designers notice problems both in their algorithms as described abstractly and in pictures they draw to help them design. While executing the proposed algorithm difficulties (missing steps, inconsistencies between parts of the algorithm) may arise, leading the designer to further refinements. Thus, we say that the designer's refinement process is *difficulty driven*.

[D1.5] In DVs algorithm a difficulty arose when the test involving line segments was combined with the generator of points and D1 had to modify the algorithm to accommodate this.

Here one assertion propagated by the execution process (that a point is produced by generating over the input set) contradicts another assertion (that a line segment should be the input to the test rather than the point it is handed).

[D1.6] D1 eventually changed the kernel idea from generate and test to a greedy algorithm that attempted to generate the hull points in the order they occurred on the hull polygon, using backup to handle guessing failures.

The algorithm execution also can expose opportunities for improvement or modification of the algorithm.

[D2.5] After working on a sample problem, D2 realized that the merge step would be easier if the two subsolutions shared a common point and went back and modified the divide step to ensure that that would happen.

Most people draw *example figures* during algorithm design. The examples are used initially for understanding the problem, and for reasoning about the task domain (using visual reasoning in the geometric domain) as well to help try out the partially developed algorithms in test-case execution. Often, the designers notice things about the sample figures that they were not looking for. When what the designers notice turns out to be useful in developing their algorithm, we say that they have made a *discovery*.

[D1.7] In looking at Figure 3 D1 realized that if a line segment had points on both sides of it, that segment could not be on the convex hull. D1 was executing an algorithm with a test for points being on the hull or not; the line segment in the figure was recording the fact that the points A and B had been generated so far

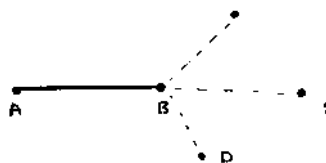


Figure 3. D1's discovery of a segment-on-hull test.

[D2.6] D2 created Figure 4 in attempting to find a merge step by considering all segments between vertices of the two hulls and testing which were in the merged hull. D2 knew that this brute force search would be too expensive, but had no other ideas. The picture reminded D2 of another unrelated algorithm (the traveling sales representative) in which a shorter path replaced two adjacent segments. D2 then applied a similar idea to the merge step, replacing segments ad and d-e by segment a-e (D2's picture was not actually labeled). The generalization D2 made was that convex angles in the merged hulls were to be replaced by a segment connecting the two end points.

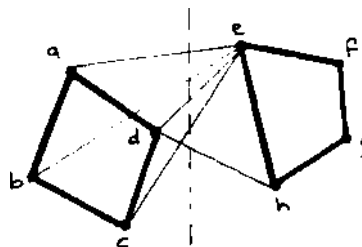


Figure 4. D2's discovery of a merge operation

Do you have enough information yet to guess the algorithms? (The discoveries are described in more detail in cite(Kant82b).)

Some other observations the designers made would have allowed only small optimizations.

[D1.8] D1 noticed that points are always on the same side of the (directed) line segments of the hull.

While discovery is not a voluntary process that can be planned as a design step, it does arise from the process of making observations. The discoveries in our study all occurred when the designer was looking at a sample figure created for one reason and recognized a geometric property, or key step from another algorithm, that would solve an outstanding goal. That goal was not the one the designer was currently worrying about (finding a test for a point being on the hull, finding a way to tell if a segment was on the merged hull), but it was usually not completely unrelated (finding a segment test rather than a point test: finding a different type of merge step). Thus, discovery could be characterized as serendipitously satisfied goals

Both key observations in the problem domain and knowledge of design principles are usually necessary for clever design. Most algorithms published in papers, or given as exam problems have at least one good observation or trick that is novel at the time of the design otherwise we would probably say the algorithm is "obvious" or is "just" a brute force algorithm. Each of the tricks, must be stumbled upon as a discovery unless it is already known to the designer. Good tricks are eventually refined into principles but everything is a trick the first time each designer encounters it.

Although there is an element of chance in the discoveries there is no lack of readiness on the part of the designer. The designer can be prepared both with immediate goals to exploit the observations and with a good understanding of design principles to fit the discovery into an overall algorithm. The more experienced and disciplined the designers, the better prepared they are for the discovery. An "experienced" designer is one with knowledge not only about algorithm design but also about problem domains. Domain knowledge can be derived either from past attempts at the problem or from experience with similar algorithms and domains (or different domains but the ability to reformulate problems in terms of other domains).

#### E. Verifying Correctness

Our designers determined whether their algorithms were correct primarily by testing them on specific examples and observing whether there were any difficulties. Symbolic execution can in fact be made to do the job of full formal verification. To do this, the algorithm is executed on symbolic objects and all assertions are propagated to determine whether the results of the algorithm (and its subparts) match the specifications. If a specification includes performance constraints, then verification must also include an evaluation (see Section IIF) to determine whether the solution is efficient enough (in time or space complexity) according to the expectations.

During the initial algorithm design, the designers ignored "details" such as base cases or initialisations, boundary conditions degenerate inputs or unresolved notes to themselves, but they were more careful about this if they were attempting to determine if the algorithm was complete or correct.

[D1 9] When D1 was asked for an algorithm summary during a pause, the response was that it wasn't an algorithm yet because the case of the first point not being on the hull had not yet been tested.

The heuristic is to get an algorithm for the general case first then worry later about modifying it to take the exceptions into account. Although some methodologies claim to eliminate the concern with special cases (for example, [11]) they require that the specification or invariant be precisely stated before design begins. This is often difficult to accomplish. For more complex algorithms handling the exceptions can itself require a major problem solving activity and may yield new insights into the problem or solution.

#### F. Evaluating Plans, Refinements, and Solutions

The descriptions of the processes used in design did not detail how plans, refinement steps, and overall solutions are evaluated. Evaluation can be based on specific knowledge about the algorithm design principles being applied or on an analysis of the cost of the algorithm and its subparts.

If the designer has the appropriate rules about the algorithm design principle and the domain, then the refinement process can be smooth and top down. For instance the appropriateness of the kernel ideas selected by the designers depends on the quality of their knowledge of algorithm design principles. One can really observe here what expert systems researchers call domain specific knowledge. Generate and test is usually the fall back idea, which is sometimes very efficient (linear in the input size) and sometimes not. After an algorithm based on a kernel approach was sketched out, or after the approach seemed to be failing, some designers went on to an alternative approach.

[D1.10] After completing the revised algorithm for generating segments and testing whether they were on the hull, D1 determined that the run time of the algorithm was proportional to the cube of the number of input points. Declaring that this algorithm was only a "first shot" D1 went on to consider a dynamic programming approach and eventually to try divide and conquer.

(D4 1] In another problem involving finding intersections of line segments another designer D4, noted that there was a straightforward approach having to do with considering all pairs of segments, which was  $N$  squared. However, D4 felt that there ought to be some way to use sorting in the solution to get an  $N \log N$  algorithm.

When experts (people with a strong background in algorithms and in the subject matter of the problem) design, they consider a variety of alternative refinements, select the best (remembering the rest for possible later use), and apply it to advance the design with one more level of detail in the refinement process. What is "best" is based on efficiency in the cases of algorithm design we studied, but is based on ease of implementation or modification in other cases. In expert design, the breadth first process tends to be followed for all aspects of the design at a given level, with interactions between the different parts of the design predicted and taken into account.

In contrast, if the designer's only idea is naive (use sorting somehow), then the technique of executing hypothesized algorithm parts is more likely to be followed in a depth first search from which the designer may never successfully return. (The idea may not have been wrong, but the designer may not have had the knowledge to carry it through.) Experts as well as novices are prone to a satisficing style of design when they are under pressure and don't have time for more exploratory design. Of course they are better at it since they have more experience, can make better predictions, and guess right more often.

Even when performance constraints are not explicitly specified the designer often evaluates an algorithm or algorithm stop's performance relative to other alternatives or to known or estimated lower bounds. Extreme cases of inputs may be tested to estimate worst case performance. Complexity analysis may be earned out in parallel with execution and verification by more experienced designers, or may be an explicit subtask of a conscious evaluation.

[D2.7] After discovering the way to merge by removing convex angles, D2 estimated the run time of the divide and conquer algorithm by arguing that even for the worst possible input, the merge time was linear in the number of points on the two subhulls and therefore the overall run time was acceptable.

[D5.1] Although D2 did not draw an additional figure to analyze the worst case, both D1 and another designer D5 did. Even though their algorithms differed from that of D2 and from each other, they were both concerned with the same potential problem and drew similar pictures (see Figure 5)



Figure 5. A worst case input for divide and conquer

D2's final algorithm had a prepass step to sort all the points according to their X coordinate. The basic algorithm was to divide the input through the point closest to the median, recursively find the convex hulls of the two resulting point sets, and merge the solutions back together by eliminating concave angles (starting from the shared point) 1 The base case is that two or three point input sets can be made into convex hulls immediately.

#### G. Control Issues

The design processes described in the preceding sections do not always run to completion and do not take place in any fixed order. Evaluations within each step, as described above, may cause the designer to terminate one approach and go on to another. The ordering of the design processes (including when they begin and end) seems to arise naturally out of the mechanism of trial execution.

Selecting a problem to work on is a natural consequence of the problems exposed by symbolic and test-case execution. The character of the elaboration process appears to be an progressive deepening that takes each of the constructs in the algorithm a little further, sometimes backing up to higher levels to keep the overall picture in mind. However, the development of the different parts of the algorithm is not always even. If one aspect of the algorithm is a potential problem (i.e., other parts of the design depend on it and the outcome is uncertain), then it is more likely to be expanded to ensure that the algorithm as a whole is feasible. If it has an obvious solution or refinement and the implications of that decision seem well understood, at least at the current level of detail, it is not considered further. (Of course the assumptions may be wrong.) New components of the design are refined in the order they are executed, subject to the two previous considerations.

Verification and complexity analysis also seem to be achieved in part by propagating assertions during execution. Thus, other processes that contribute to control fit in nicely with this basic mechanism and can occur at the same time.

In short, design processes are applied as appropriate. Control is not a special source of intelligence. It comes out of responding to the data and out of the problems and opportunities arising during execution.

### III. Locating the Problem-Solving Power

An important question to ask about any agent that exhibits intelligent behavior is where the knowledge and problem solving power lie. Knowing the loci of intelligence gives us some clues for how to produce similar behavior automatically. Thus, we have

attempted to formalize the problem-solving behavior we observed in our designers in terms of concepts that lend themselves to automation

One common view of problem-solving behavior is that it is basically search in a problem space, with knowledge used to limit search. Knowledge is carried by the problem spaces themselves, in what objects and operators they have available and in the heuristics they have for when and how to apply the operators. In this view, problem solving is a process of repeatedly changing a *context* by selecting a *goal* to achieve, a *problem space* to work in to attack that goal, a *state* within that space to work on, and an *operator* (and instantiations of its *arguments*) to transform the state [18]. Different types of knowledge can be identified with the selection process for each element of the context

#### A. The Power of Search

In design as in most tasks requiring intelligence, both search and knowledge are needed. Search is the backup for missing knowledge and can never be completely eliminated. It can take place at the very high level, such as searching for a kernel idea for an algorithm or at the very low level, such as deciding how to instantiate an operator argument. Although at any level knowledge limits search when possible and gives clues about how to explore the problem spaces in a reasonable way, the ability to search is, in itself a source of power.

In design, for example, search permits the creation of algorithms by trial and error in the absence of complete knowledge. Algorithm components can simply be hypothesized and then the algorithm as a whole tested to see if it satisfies its specifications. If only the objects and operators that formally specify and manipulate algorithm descriptions are available (i.e., there is no other model of the problem domain), then designing an algorithm requires the use of formal definitions of the concepts used in the problem specification and, recursively, of its subcomponents. However, more power than this is available to human designers and can be made available for automated design through the use of multiple problem spaces.

#### B. The Power of Multiple Problem Spaces

From our observations we conclude that each designer works in several different problem spaces during design (similar observations are described for other tasks in [20]). The details of the problem spaces differ from designer to designer, but there is a remarkable consistency in the types of problem spaces used.

We observed our designers working in four spaces, two of which are extensions of another space. The two main spaces were (1) an *algorithm design space* that carries the knowledge of what is achievable in standard computer systems and of domain independent algorithm design principles, and (2) an *application domain space*, such as one for geometric and visual reasoning (The algorithm design space is also a domain space relative to design as a whole.) The two extension spaces have the same objects as the first two spaces plus additional objects and/or different sets of operators. (3) An *algorithm execution space* is an extension of the algorithm design space that has as new objects data items that carry information in the form of assertions about their execution history and has new operators that execute components in the design. (4) An *example generation space* is an augmentation of a task domain space in which figures are marked as standard examples, degenerate cases, counter examples and the like, and in which there are new operators to produce the examples.

The necessity for different problem spaces is a result of the requirements of different types of knowledge. For example, what is possible or efficient in the domain (problem space) of algorithms for conventional digital computers is sometimes quite different from the way people reason visually or from what can be done with analog devices. (Consider solving the convex hull problem by pounding nails into a board to represent the input points and then stretching a rubber band over the nails and letting go.) The problem spaces that express such knowledge differ in the objects and operators included, the properties of objects or relationships between objects, and heuristics for how to control the applications of operators.

Having knowledge represented in a domain space as well as in an algorithm space gives the designer the power to create algorithms even in the absence of formal axioms about specification concepts such as polygon containment. The problem can be solved by generating constructs in the algorithm space and testing the proposed algorithms on examples to see if they work. This technique relies on the ability to generate examples to use as test cases. Example generation depends on knowledge of the domain space as well as knowledge of the goals in the algorithm space (say to determine whether a typical or degenerate example is desired) If a domain concept is not formally axiomatized, the designer cannot do any formal symbolic reasoning such as full verification or correctness-preserving derivation. However, by making some conjectures about the domain and validating them with test case execution, the designer can reason formally about the rest of the algorithm

[D2.8] Having knowledge from the domain space of what line segments were on the merged hull allowed D2 the hope of finding an operation that would test where proposed segments were correct.

For each of the problem spaces relevant to design, we can ask what knowledge is available for recognizing when context elements should change: how does a system recognize when goals are satisfied or when new goals should be attempted, when the problem space should be changed to work on the different type of goals, what state to expand within a problem space, and what operator to apply and how to instantiate the operator. Examples of the different types of knowledge contained in problem spaces will be given in Sections III.F through III.I

First, some aspects of problem-solving power that cut across problem spaces are discussed. This power can be cast as knowledge that allows the designer to avoid search.

### C. Knowledge in Recognition

The ability to recognize objects and to recognize the applicability of operators is a major source of power in problem solving. The search process is not driven by an algorithm that selects context elements in a fixed order but rather by recognition rules that observe when some context element should change: for example, when a goal has been satisfied or when an operator would help change state in a desired way. The conditions for recognition can be symbols in the algorithm design space or visual images from the domain space. These clues can involve goals, points of view or other objects in the problem-solving context whose inclusion as a clue was only accidental to the formation of the recognition rule. A very large number of recognition rules may be present. However, the conditions that are monitored must be computationally simple, involving only straightforward matching.

An example of the role of recognition is its use in discovery, a key process in algorithm design. Discovery depends on generating examples to work with and then noticing properties about them or reasoning about them. The recognition processes usually take place in the domain space, but what is noticed depends on the goals of the problem solving (and the content of the recognition knowledge).

Recognition is also important in example generation, which is constrained by the goals of the problem solving (is it to be an average case, degenerate case, initial or base value, counter example, used by efficiency analysis, etc.), but depends on knowledge of the domain and recognition of successful construction of the example in terms of domain properties.

[D1.11] D1 first generated points A, C, D, and E in Figure 2 as an initial test-case example but then noticed that the example was degenerate since all points were on the hull and added a fifth point (B) in the center to remedy the difficulty. The points were not labeled at that time.

Non-symbolic recognition and processing (such as visual reasoning) is clearly important in designing computational geometry algorithms, but is it really important in all domains, such as that of algebraic problem solving? At least for some people, it is. Built in visual operators are better at some types of processing and provide another perspective on a problem. They may suggest approximations or fortuitously counterpose objects that would not be related by a general symbolic reasoning process.

### D. Knowledge in Execution

Trial execution in algorithm design serves the purpose of controlling the order of the refinement process (see Section II.C) and limiting the inferences made as well as its more common functions of debugging and verification (see Section II.E and [6]).

The nature of creative algorithm design requires some mechanism for inference, whether it is a full theorem prover small set of simplification rules, or something in between. Making all possible inferences during algorithm design would be very expensive computationally. Execution is a way to focus attention on certain assertions in the algorithm description space and certain parts of pictures in the domain space so that inference and recognition only have to take place over a smaller set. The execution techniques limit the inferences and constraint propagations to those most likely to be useful for the current stage of the design. Avoiding the extensive search of theorem proving or uncontrolled inferencing through execution is a form of knowledge about design. This topic is discussed more thoroughly in other papers [26], [7].

### E. Knowledge in Efficiency Information

Efficiency knowledge serves as an evaluation function throughout the algorithm design process, not just as an evaluation of complete designs. Information about potential run time or space use serves as a rough guideline in the selection of a kernel idea and during refinement (D2 knew that the merge step had to be linear to get the desired overall performance [D2.6]\*) as well as after an algorithm sketch is complete (D1 decided that cubic performance was probably not the best possible [D1.10]).

\*In the remainder of this paper, labels following descriptions of bits of knowledge refer back to parts of the design story where they are used

Efficiency knowledge can take many forms, including assertions about the run time of specific operations or algorithms, assertions about the intrinsic complexity of problems, rules for how to analyze algorithms, and rules for setting constraints on what performance must be reached on a subpart of an algorithm to guarantee overall performance.

Efficiency knowledge is generally contained in the algorithm design space. In fact, some heuristics in the algorithmic problem space depend on the cost model for the target architecture, efficiency knowledge can be applied in selecting a plan [D1.1], in evaluating refinements for steps of the algorithm [02.7], or in evaluating the algorithm as a whole [D1.10].

## F. The Algorithm Design Space

In algorithm design, it is sometimes difficult to come up with any reasonably effective solution,\* although some problems have simple brute force solutions. (Consider the problem of finding the closest pair of points in a point set. You can probably see a simple algorithm for solving the problem immediately.) Since algorithm design involves searching in a space not dense in solutions, dead ends are a serious problem, and knowledge of what design principles and domain facts are relevant is almost a necessity (as is the ability to reason and recognize in other spaces). Such knowledge can help decompose the problem or select and instantiate operators in the problem space.

Designers have variants of the algorithm design space that depend on their assumptions about the target architecture as well as on their overall knowledge of design principles. If the algorithms were to be programmed on an architecture with pipelined or distributed processing or associative retrieval, the representations for algorithms and heuristics for how to design might be greatly different. Some designers make (at least implicit) assumptions about the target architecture from the beginning of algorithm design, although it is preferable to stay independent of the target as long as possible.

The knowledge in the algorithm description space includes facts about mathematics, logic, arithmetic, or algorithm design principles. The knowledge can be in the form of both object descriptions and operators on those objects. Other knowledge can be represented by rules about when to change the problem solving context.

### 1. Objects and operators

The basic objects for describing algorithms in the algorithm design space are *components* that specify basic types of processing. These components may test whether a property holds, generate the elements of a set one at a time, achieve an input/output relationship, apply a domain operator, select a subpart of a compound object, or modify a memory of objects.

The algorithm components are connected by *links* that allow flows of data and/or control and may be augmented with *assertions* about their properties or about their relationship to other objects or operators in any of the problem spaces. For example, a selection criterion might be to pick the bottom left point from a set of points. New components can be defined in terms of old ones by adding additional standard inputs or outputs or by adding assertions, or a component can be defined as configuration of other components.

The assertions associated with components may include information about the types of data objects expected as inputs or outputs or other preconditions or postconditions of processing, the ordering constraints on a generator, the criteria for selection, the initialization of a memory, expectations or conclusions about the time complexity of the algorithm (component), constraints on the order of execution of the algorithm components, notes about the algorithm (such as it has not yet been tested for the initial point lying inside the hull).

Since algorithms usually manipulate some sort of data, there are also representations of the common mathematical concepts such as numbers or symbols and of sequences or sets of other objects. Assertions about these objects can be attached to descriptions of the object type or to *item*-, that represent specific data.

*The number of combinations of pairs, from a set of elements n;*  
*for partitioning into N - injured (D4.1)*

*Divide and conquer algorithms often have run time of*  
*N log N. [D2.10.11]*

The operators in the algorithm description space are simple (syntactic) editing operations that add or modify components, links between components, and assertions. The knowledge is all located in the rules that suggest instantiations of the type of components to create, the specific components to link, and the details of the assertions to be added.

### 2. Operator selection

Selecting an operator (and instantiating it by selecting values for its arguments) can be made more effective through the use of knowledge about general algorithm design principles and about algorithms in a particular domain of application. This knowledge will be expressed here as rules. Other such knowledge, such as how to handle specific problems raised during execution (the equivalent of a difference table for means ends analysis) also limits the amount of search necessary for operator selection.

The following set of rules about operator selection and instantiation is merely a representative sample of the knowledge that an algorithm designer (human or otherwise) might have (not every designer has the same knowledge, of course). Many other rules would add their suggestions and vetoes about what to do. If there is no consensus about what operator to apply, the fall back is search through the suggested possibilities.

*If a component needs to be refined and its output is a subset of its input, refine the component to an element generate and test algorithm. [D1.1]*

*If a component needs to be refined and its output is a structure that must satisfy certain constraints, refine it to an algorithm that builds a minimal structure and then adds units of structure until the constraints are satisfied. [D1.6] (An instance of this rule is suggested in [3].)*

*If an algorithm looks at part of the input many times to do the same kinds of tests, try saving information rather than recomputing, say with dynamic programming. [D1.10]*

*If the characteristics of subproblems produced by the divide step of a divide and conquer algorithm are unknown, then add the assertion that they are two equal sized subproblems*

*If the characteristics of subproblems produced by the divide step of a divide and conquer algorithm are unknown, and if the set being divided is a set of points in two dimensions, then refine the divide step to be a sort of the points and a division into the points on either side of the median a line through the median. [D2.3] This has a bit of domain specific knowledge although it is in the algorithm space*

\*See Section ME or [2] for a comparison with the search problem in program synthesis

*If a component is missing a link to a required input, look for a component that has an output with the same type (or having that type as a subpart or superpart) and connect the two components.*

### 3. Changing state

The state in the context of a problem space changes primarily as a direct result of the successful application of an operator that modifies the algorithm description. If the operator application fails, and if there were competing suggestions about what operator to apply, then alternative operators still apply and another will be tried. In addition to either failing or succeeding, an operator may return a difficulty or opportunity. This becomes another goal to be worked on, perhaps in a different problem space. After processing of the new goal is complete (which may change the state in the algorithm description problem space), the rules that caused the original operator to be selected may or may not be retriggered. If they are, the operator application can be retried.

### 4. Changing problem spaces

One of the benefits of having multiple problem spaces is the ability to reduce search by working on the same goal in a different space. Some examples of rules that can cause space changes are:

*If a component needs to be refined, and its output is a construct in space X, create examples of it and notice their properties. If this rule is applied, it will cause a transfer first to the example generation space and then to the domain space X. [D1.2]*

*If a configuration of components has not been shown to achieve the specifications of the component of which it is a refinement, then symbolically execute it. [D1.9]*

*If a configuration of components has not been shown to achieve the specifications of the component it is a refinement of, and if symbolic execution has already been tried or is known in advance to be too complex to be informative, then execute the configuration on a concrete example. [D1.2]*

### 5. Goal satisfaction and creation

Recognition of when goals have been achieved, or nearly achieved, of when to give up on a goal and declare failure, of when to create new goals, and so on is crucial to enabling discoveries. Strict enforcement of hierarchical subgoaling would not allow the same flexibility and creativity. Goal change knowledge can also serve as design heuristics. Some rules that express this knowledge are:

*If an exponential algorithm is created, try to improve it or find an alternative unless it can be shown that the problem is itself exponential. [D2.6]*

*If all objects added to a set have a common assertion, hypothesize that that property holds for all elements in the set and try to substantiate the hypothesis,*

*If a component is defined by assertions that are appropriate for the level of detail currently desired (however that is determined<sup>1</sup>), then consider the component acceptable.*

*If a component is not considered to be refined to an acceptable level of detail, then create a goal to refine it.*

### G. The Application Domain Space

Algorithm designers need knowledge about their task domain as well as about algorithm design in general. As an example of a problem space describing a task domain, consider the knowledge about geometry that can be used in solving the convex hull problems.

Objects that are manipulatable in the geometric domain include points, lines, segments, angles, and polygons. Special properties of object types or of specific objects may also be recorded. For example, the degenerate case of the object type polygon could be a point or line-segment, and a triangle would be the boundary case. For a specific geometric object, properties would include being convex or being above or below a line.

The operators in the geometric domain include accomplishing such functions as drawing a line segment between two points and recognizing that a polygon is convex.

Any symbolic descriptions of the objects in a figure and assertions about the objects or their relationships are available to the other spaces. For example, in the algorithm space, assertions may serve as test predicates, comparison or ordering relationships, or criteria for extraction from compound object. Operators are available for execution, say to build a polygon in the example generation space or as an operator applied by a component in the algorithm space that can be run during test-case execution, but their internal workings are not available.

The domain space also includes recognition knowledge, expressed here in the form of rules, that if applied to a figure in the current focus of attention may cause recognition and/or the construction of a new object just as an operator application might. For example,

*If two line segments share a common endpoint, perceive the figure defined by that pair of segments as an angle. [D2.6]*

### H. The Execution Space

The problem space in which execution occurs is an augmentation of the algorithm description space. It uses the object type *item* to represent the data processed by the algorithm that flow over the links between components. The items can represent either specific objects from the domain space (point A) or symbolic objects ("a point"). Items can be augmented by properties that are known to be true of them at a given point in the algorithm execution history — that a point is known to be on or off the hull or that it is the one most recently added to a memory.

The operators in this space control the sequencing of component execution and carry out component execution. If assertions needed to carry out the operators are missing, a difficulty is returned and a new goal to handle the difficulty is created.

Some instances of rules that suggest new goals to work on are:

*If the input for test-case execution is uninstantiated, set up a goal to get an example input. This will cause a transfer to the example generation space. A particular point set would be an example for the convex hull problem. [D1.3]*

*If test-case execution shows that applying some operation will make progress toward a solution of the problem but not solve it completely, try modifying the description in the algorithm design space to apply the operation repeatedly (inside a loop).*

### I. The Example Generation Space

The example generation space is also an augmentation of another space, the domain space. Objects must be augmented by properties that describe their typical instances, degenerate instances, boundary cases, and so on, if such information is not already present in the domain space. For instance, sequences consisting of repeated copies of the same element are not typical. Some sample operators are those that add and remove elements from examples. Some sample rules are

*If creating an input to a generate-and-test algorithm and all elements in the input satisfy the test, then add another element [D1.10]*



*// creating an example for test case execution of an algorithm that has not yet been checked for correctness, pick non degenerate objects and constructors*

#### IV. Design Automation Strategies

This section summarizes the model of human design and compares it to some of the other approaches suggested for fully or partially automating algorithm design and for automatic programming. It also discusses how the methods might be extended to handle the problems in other contexts, such as interactive design.

##### A. Summary of Human Design

Several of our designers succeeded in creating convex hull algorithms. The algorithms and key discoveries of designers D1 and D? have already been described. DVs generate and test algorithm had a disappointing worst case run time proportional to the cube of the number of input points. But D1 would never have been able to design the anticipated linear algorithm. It can be shown that the problem of finding a convex hull is related to the problem of sorting, so under conventional assumptions it must be an NlogV problem. Eventually O1 went on to try a divide and conquer approach that, with a little help from the experimenters became a successful NlogV algorithm similar to D2's. Some other designers successfully recreated some convex hull algorithms that they had heard or read about but did not remember very clearly. (Many interesting convex hull algorithms have been described in the literature [19].) Still other designers failed to find any algorithm at all. We also gave our designers some other problems. They were asked for algorithms to find the closed pair of points from a given set or the intersection points of a set of vertical and horizontal lines. Most designers quickly suggested brute force algorithms (which have a worst case run time that is the square of the size of the input) but were unable to find any of the faster algorithms.

The methods observed in human design are quite varied. Selecting and sticking with a kernel idea provides a necessary focusing of attention, and using execution as an assertion propagation mechanism continues that focus and avoids the extensive search process that unlimited inference or search through the network of all refinements would entail. Of course if specific knowledge about the domain or algorithm design is available, it can be used to limit search by suggesting refinements directly. A powerful source of creativity is the use of visual reasoning about specific examples, which paves the way for discoveries about key concepts in algorithms. Although our current set of studies of human designers has provided many good ideas for a model of design, we would like to do more studies on other types of algorithms and on even more expert algorithm designers.

In general, the designers' success was highly correlated with their interest in and background in algorithm design. Some problems that they had stemmed from an incomplete (or totally absent) understanding of design principles such as divide and conquer (which is very relevant to the examples we gave). Other problems seemed to be due to impatience with methodically following a design strategy. In some cases, the designers tried to mix aspects of the design from two different approaches. This typically failed when they tried to mix subparts of different types of principles but succeeded when they tried to reuse facts or theories from the geometric domain that were learned in an earlier design.

##### B. Automatic Programming

Automatic programming is that ever receding goal of automating the programming of everything the user wants with a minimal amount of specification. Automatic programming encompasses (1) algorithm design, (2) program synthesis, and (3) the problem of managing complexity in programming in the large. Algorithm design has been defined in Section IA as the process of producing a computationally feasible program sketch (that is relatively complete and consistent) from a specification of what is to be accomplished. We refer here to the hour level form of algorithm design, not research design. This routine design often precedes program synthesis. Program synthesis is the process of choosing data structures and access functions to transform a given algorithm specification into concrete code in a conventional programming language. Like algorithm design, program synthesis requires intelligence, especially to produce extremely efficient code, but it probably can be achieved with more straightforward techniques.

As has been pointed out by others [2-10], full fledged automatic programming requires the incorporation of domain knowledge as well as detailed coding knowledge. Furthermore, programming in the large must be supported by effective bookkeeping. There are few concrete results in this area, however. The notion of working in multiple spaces, and in a domain space in particular, may prove valuable in automating the entire programming process.

##### C. Formal Derivation

The formal derivation approach has been proposed for both algorithm design and program synthesis [25, 5, 21]. Formal derivation methods share with the design methods described here a refinement strategy based on a few, largely syntactic, transformations, but differ in that the transformations preserve correctness. It is assumed that the specifications are correct and complete, and since the transformations require and guarantee correctness, then the intermediate states and the result are also correct and internally consistent. The operations of the transformations — defining new constructs, expanding definitions ("unfolding"), noticing instances of definitions that have arisen after rearrangement and simplification of the algorithm constituents ("folding") — are similar to the processes that we have noted in human design.

One way that the formal approach differs from the model of design described here is that it requires that terms be defined by axioms or equations and does not allow the use of terms defined only in a domain space. Also, in the formal approach, transformations are instantiated via axioms about the domain or algorithmic constructs, in the model of design described here, they can be instantiated by similar knowledge based on formal definitions, by arbitrary selection, or by guesses based on observations of the domain. As discussed earlier, people can sometimes derive algorithms even if they do not have formal definitions of all the concepts. They need only have operators in the domain space that recognize the concepts, more primitive operators in the domain space that can construct the structures they want to recognize and techniques for implementing the constructive operators in the algorithm space. In contrast, the formal derivation approaches often have problems with controlling the search process and with creating useful auxiliary definitions — the "aha" or "eureka" steps are often definitions inserted by human interaction. These problems result from there being no clues in the formal approach about how to introduce the right interesting knowledge.

Another way the formal approach, with its requirement for consistency and completeness, differs from human design is in the handling of boundary conditions and base cases. The formal

approach requires that these be defined early on, almost the opposite of the human approach. Getting the details of the boundary conditions right is one cause of the search problem in formal systems — there are many ways to define these conditions, and selecting the precise specifications or introducing conditionals and filling out the details adds complexity.

For some people, the discipline of taking care of details with a standard methodology releases their creativity. On the other hand, many people find it difficult to state invariants precisely if they must be absolutely correct. Getting the main idea of the invariant is crucial to solving the problem, but stating it formally to avoid such problems as fencepost errors makes it tedious and not obviously productive. For these people, getting the details right immediately is extremely difficult; the overhead of internalizing this methodology is prohibitively high.

Formal derivation systems are being augmented with more detailed knowledge about design techniques so that the search control can be more goal oriented [9] and also with knowledge about example generation [4]. However, this still doesn't postpone settling all the details (having a domain space lets you finesse formalizing them) or say where the creative definitions come from (cross fertilization from domain spaces and other algorithms).

#### D Inductive Inference

Inductive inference from examples is another technique that has been explored, but more for the construction of small programs than for the design of algorithms or large systems. Unambiguously specifying the input/output behavior of algorithms with examples is easier than so specifying the behavior of large programs. However, the inductive approaches usually rely on problem solving using a small set of schema, with little ability to improvise if none of the schema match. If the target language is logical equation based language with a search mechanism built into the interpreter, then this approach may work [24]. But it is unlikely to produce clever algorithms in conventional languages.

#### E. Program Synthesis by Refinement

The program synthesis problem is complementary to that of algorithm design, although we would expect that many of the same problem solving techniques are used. The stage at which the algorithm design process stops — when an algorithm is "understood" — should provide an appropriate specification or starting place for program synthesis.

The standard refinement paradigm in program synthesis [13, 23] is to apply knowledge-based rules and search over that knowledge; no creativity is introduced. The search problem is a bit different since once an algorithm is well defined, the program synthesis problem is usually to find a more detailed program in a standard programming language selecting concrete data structures and accessing operations. Usually the search space is dense in correct solutions that vary in efficiency, reliability, modifiability, and so on [2]. Past research has investigated the control of the search by efficiency (for example, [14]). Such control is not a definitive solution, but many approaches have been prototyped fairly successfully.

As in most expert systems, it is assumed that all the knowledge about how to refine programming constructs is present in the refinement rules. In contrast, the hypothesize and test technique in the design model presented here allows the discovery of new programming techniques. The price paid, of course, is that more search at the lower levels is required, and this search is not as easily controlled by efficiency rules.

#### F. Program Synthesis by Design

We hope that algorithm design research will result in aids for program synthesis that avoid hand coding of all the refinement rules. The initial knowledge base requirements should be simplified considerably as a result of the more generic problem-solving abilities such as trial execution, with its low-level means ends analysis and search, and domain space reasoning. Putting in more of this creativity should make the automatic programming process more flexible and robust and may even produce better programs.

#### G. Interactive Tools

An interesting question to ask is, does this knowledge suggest any other tools to aid in the design process? Are there some interactive tools that might help people in the design process? Or is there some novel mix of human and machine power that could lead to even better design?

The conventional wisdom is that people have better insight and machines are better at the details. Following this wisdom, the machine could suggest the full range of possible approaches at any one step and the person could decide which to follow, providing the search control.

We could augment this plan by observing that execution is a powerful technique in design. Programs are good at methodically following algorithms for execution, but people frequently see what they expect and miss some of the problems. This would suggest machine support for execution of designs. The execution would expose problems and inconsistencies that people might skip over and the people could suggest some solutions to the problem or suggest new directions to follow.

In addition, the machine support could include a set of rules that continuously monitor simple features of the design, providing a check that preserves almost-correctness but does not guarantee a complete validation. In effect, this makes the machine a sounding board for human design, just as colleagues act as sounding boards. People explain their ideas to colleagues so that they are forced to look at their design from other perspectives (with different assumptions) and go through the design one more time in explaining it. A machine might serve the role of a colleague.

Building the human/machine communication interface is the hard part of following through with these plans. The two agents must speak the same language and each must be able to track what the other is doing. This may turn out to be even harder than full automation.

#### H. Other Design Tasks

There are a variety of other design tasks, such as engineering design or VLSI design. Although each of these tasks has its own unique characteristics, we may hope that some of the concepts discussed in this paper may be relevant to these tasks.

## V. Conclusions

The essence of the model of design presented here lies in its informality and its use of multiple problem spaces, including example generation and trial execution based on both the domain space and an algorithm design space. These techniques provide a focus of attention to limit search and enable the discovery of key concepts. The model shares problem-solving techniques with many of the other approaches, but rather than having a single monolithic plan of attack, it shifts techniques depending on the knowledge available.

Several areas need further formalizing and testing. The models of the processes of discovery and visual reasoning must be extended. Learning and database issues should be explored further. For example, what are the appropriate organization and retrieval techniques for large amounts of information so that key ideas in algorithms and derivations are accessible when relevant? Being able to learn automatically depends on appropriate accessing and on general problem solving techniques,

The interactions between search, domain knowledge, and programming knowledge seem important in tasks of any appreciable difficulty, including automatic programming and the next generation of expert systems, but several questions about these interactions are still unresolved. For example, it is not well understood how to determine when to stop refining at a given level, how problem spaces are created from problems descriptions, and so on.

Understanding the design process impacts other branches of AI. Those that include design tasks, discovery, visual reasoning, the use of examples, and interaction between different types of knowledge could be compared to algorithm design in their organization of knowledge and use of problem solving techniques. Answering the questions posed for design should shed some light on the general issues in other domains. A side effect of automation, the formalization of algorithm design, analysis, and optimization principles, could also be useful in teaching. Our observations of human design show that examples are useful in the absence of knowledge and therefore probably necessary to teach the knowledge, but having explicit principles is more efficient for the designer.

In summary, the model of design presented here is a good start on understanding algorithm design. The attempt to formalize the model lays a substantial part of the foundation for automation.

#### Acknowledgments

The research described here is joint work with Allen Newell and David Steier. Mary Anne Wolf recorded and transcribed many of the protocols. Thanks to Jon Bentley for helping to instigate these studies. Thanks are also due to the designers who participated in the experiments for their interest and cooperation; they shall remain anonymous for obvious reasons. Programming support for the implementation and additional input has been provided by Brigham Bell, Lisa Covi, Billy Kim, Roland Kovacs, Deepak Kulkarny, David Marshall, Jim Muller, Ed Pervin, Eric Schwabe, Mark Taylor, and Ross Thompson as well as Wolf and Steier. Allen Newell, David Barstow, David Steier, and Sol Greenspan made valuable comments on earlier versions of this paper.

#### References

1. Adelson, B. and Soloway, E. A Model of Software Design. In *The Nature of Expertise*, Chi, Glaser and Farr, Eds.. Lawrence Erlbaum Associates, in preparation.
2. Barstow, DR. "A Perspective on Automatic Programming." *The AI Magazine* 5, 1 (Spring 1984)
3. Barstow, DR. The Roles of Knowledge and Deduction in Algorithm Design. In Biermann, A. W. (editor). Ed., *Automatic Program Construction Techniques*, McMillan, 1984. Chap. 10, pp. 201-222.
4. Bibel, W. and Horing, K. M. LOPS: A System Based on a Strategic Approach to Program Synthesis. Proceedings of the

- International Workshop on Program Construction. France. September, 1980.
5. Biermann, A. W. (Ed.), *Automatic Program Construction techniques*. McMillan, 1984
6. Chandrasekaran, B. and Radicchi, S (Ed). *Computer Program Testing*. North-Holland. 1981
7. Cohen, D. A Forward Inference Engine to Aid in Understanding Specifications. Proceedings of AAAI 84. 1984.
8. Ericsson, K. A. and Simon, H. A. "Verbal Reports as Data." *Psychological Review* 37, 3 (May 1980). 215-251
9. Feather, M. S. "A System for Assisting Program Transformation." *ACM Transactions on Programming Languages and Systems* 4. 1 (1982), 1-20.
10. Green, C., Luckham, D., Balzer, R., Cheatham, T., and Rich, C. Report on a Knowledge Based Software Assistant. Tech. Rept. RADC-TR-83 195, Kestrel Institute, August, 1983.
11. Cries, D. *The Science of Programming*. Springer-Verlag, 1981.
12. Jeffries, R., Turner, A. A., and Poison, P. G. The Processes Involved in Designing Software. In *Cognitive Skills and their Acquisition*, John R. Anderson, Ed., Lawrence Erlbaum Associates. 1981. ch.8
13. Kant, E. and Barstow, D. R. The Refinement Paradigm: The Interaction of Coding and Efficiency Knowledge in Program Synthesis. In *Interactive Programming Environments*, Barstow, D. R., Shrobe, H. E., and Sandewall, E., Eds., McGraw-Hill. 1984. pp.487-513.
14. Kant, E. *Efficiency in Program Synthesis*. UMI Research Press, 1981
15. Kant, E. and Newell, A. Naive Algorithm Design Techniques. A Case Study. Proceedings of the European Conference on Artificial Intelligence, Orsay, France, July, 1982
16. Kant, E. and Newell, A. "Problem Solving Techniques for the Design of Algorithms." *Information Processing and Management* 20, 12 (Spring 1984).
17. Kant, E. and Newell, A. An Automatic Algorithm Designer: An Initial Implementation. Proceedings of AAAI-83, 1983.
18. Laird, J. E. Universal Subgoaling. Tech. Rept. CMU-CS 84 129, Carnegie Mellon University, Computer Science Department, May, 1984.
19. Lee, D. T., and Preparata, F. P. "Computational Geometry: A Survey." *IEEE Transactions on Computers* C 33, 12 (December 1984)
20. Newell, A. and Simon, H. *Human Problem Solving*. Prentice Hall, 1972
21. Partsch, H. and Stembruggen, R. "Program Transformation Systems." *Computing Surveys* 15, 3 (September 1983).
22. Polya, C. *How to Solve It*. Doubleday Anchor, 1957
23. Rich, C., and Shrobe, H. Initial Report on a Lisp Programmer's Apprentice. In *Interactive Programming Environments*, Barstow, D. R., Shrobe, H. E., and Sandewall, E. Eds., McGraw-Hill. 1984. pp. 443-463
24. Shapiro, E. Y. An Algorithm that Infers Theories from Facts. Proceedings of IJCAI 81. 1981. pp. 446-451
25. Smith, D. R. Top Down Synthesis of Simple Divide and Conquer Algorithms. Tech. Rept. NPS52 82 011, Naval Postgraduate School. November. 1982.
26. Steier, D. M. and Kant, E. Symbolic Execution in Algorithm Design. Proceedings of the Ninth International Joint Conference on Artificial Intelligence, Los Angeles, CA. August, 1985