NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

Applications of AI in Design Research
at Carnegie Mellon University's EDRC

by

Art Westerberg, Ignacio Grossmann, Sarosh Talukdar,
Friedrich Prinz, Steven Fenves, Mary Lou Maher

EDRC 05-30-89
Carnegie Mellon University

APPLICATIONS OF AI IN DESIGN RESEARCH AT CARNEGIE MELLON UNIVERSITY'S EDRC

Art Westerberg, Ignacio Grossman, Sarosh Talukdar, Friedrich Prinz, Steven Fenves, Mary Lou Maher

Engineering Design Research Center, Carnegie Mellon University. Pittsburgh, PA 15213 USA

PHILOSOPHY OF THE EDRC

In 1984 the National Research Council issued a report on industrial competition in the U.S., noting that the U.S. was successful at new product innovation but was falling behind others - notably Japan and Western Europe - in bringing these products into the market place. Since 1985 the National Science Foundation has established 18 university based Engineering Research Centers as centers of excellence, each with a major goal to respond to this report. Our Engineering Design Research Center (EDRC) was started as a part of this program in May, 1986. The goal of the EDRC is to play a leadership role in developing and integrating concepts and methodologies that will allow U.S. industry to design much better products much more quickly.

Approach

In our research planning we have used both a top down and a bottom up strategy. Using the former we established three thrust areas:

1. Gaining a broad understanding of synthesis (the systematic generation of design alternatives and the selection of better ones based on incomplete information.
2. Understanding life cycle issues (how to create products which are manufacturable, disposable, maintainable, etc.).
3. Developing the concepts needed to create design systems that allow the rapid creation and delivery of new as well as existing methodologies.

fF University Libraries
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Three laboratories have been created, one for each of these thrust areas: the (1) Synthesis, (2) Design for Manufacturing and (3)Design Systems Laboratories,

In our bottom-up approach we have deliberately established projects in a variety of domains, including chemical processes, steelmaking, construction, VLSI, materials properties, and the manufacture of mechanical parts. We are developing generalizations from these projects, which in turn give us new insights into design that become part of our top-down planning.

Figure 1 illustrates some of the ideas used to aid in our original planning. It shows three levels at which one carries out a design. At the first level is the creative activity, where concepts are established on which the design will be based. This level is the least well understood and has very few computer tools to aid it. At the other extreme is the third level at which reside those activities that provide evaluations to check the performance of proposed designs. Computer based quantitative analysis tools exist here, and they certainly are the most well-developed in industrial practice. Between, at the second level, are the activities which connect concept generation and performance evaluation. Here the designer performs what are sometimes the fairly routine activities of generating alternatives, preparing different input for analysis tools, running these tools and interpreting the output, etc. At this level few computer based tools currently exist, although it seems many could.

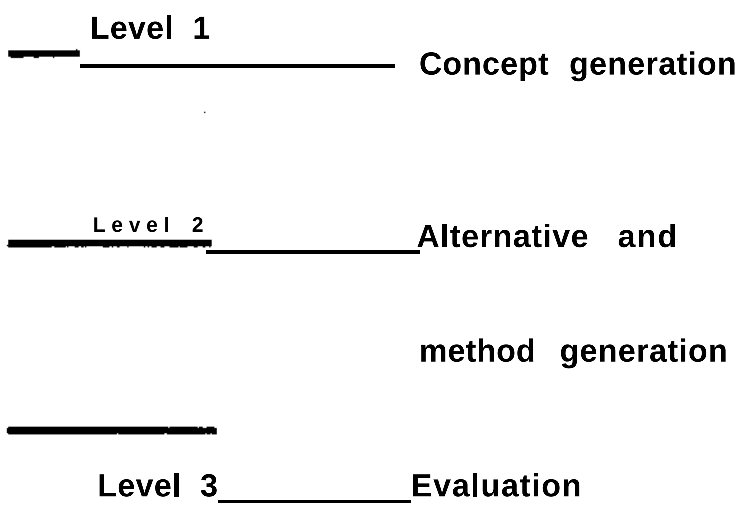

Figure 1 Levels of Activities for Design

Role of AI

We sec several areas in the different levels of design activities where the use of AI techniques play a role. For example, at the top level there could be computer aids based on AI techniques to search an existing information base for analogies to generate alternative concepts for a new design.

At the intermediate level one might employ expert systems to create and search opportunistically the space of design alternatives, using as needed available analysis packages. Finally at the third level "critics" based on knowledge can assess the quality of a proposed design against criteria for which there are as yet no quantitive models to assess performance. For example, how manufacturable is a part or how safe is a process?

With these ideas in mind, we now examine the projects within the EDRC that involve the use of AI techniques. These projects are grouped by the three thrust areas described above.

SYNTHESIS OF DESIGN ALTERNATIVES

Approach

In the area of synthesis, we are developing domain-independent synthesis methodologies to aid engineers in quickly identifying improved designs from a very large number of alternatives. The motivation behind this emphasis lies in the following facts:

1. The number of potential alternatives in the design of engineering systems is usually very large.
2. Design decisions at the synthesis stage have a great impact on the economics and quality of performance of a system.
3. In practice, only few alternatives are commonly examined due to the lack of synthesis tools.
4. The synthesis of engineering systems, which lies at the heart of the design process, is an area that is still poorly understood.

Hence, synthesis is of great importance to both the theory and practice of design, and is an area that requires substantial research effort. The synthesis of engineering systems is a very challenging area because it deals with the creation and invention of systems that can satisfy certain specifications and requirements. Since these problems tend to be open-ended and ill-defined, we have decided as a first step to sharpen the focus in this area. Specifically, we are concentrating on the systematic generation and selection of alternatives for the improved integration of engineering systems consisting of existing components.

The justification for this focus lies in the fact that the largest number of design problems in practice deal with the synthesis of these types of systems [7]. Consider as examples, the design of a building in architecture and civil engineering [8], the design of a process in chemical engineering [9,10], the design of a VLSI circuit in electrical engineering [11], or the design of a servomechanism in mechanical engineering [12]. When these problems are closely analyzed, it becomes apparent that each of them is dealing for the most part with known or existing components. What each of these areas is mainly concerned with are the following questions:

1. **Which components should be in the system?**
2. **How should these components be interconnected?**
3. **How should the parameters of the components be selected?**

Even if the above questions can be regarded as taking a somewhat restrictive view of synthesis ("routine" design), it is clear that they are not trivial. In fact, design practitioners have little guidance to tackle them, since they often have to rely on their intuition, judgment and previous knowledge. It is our objective to gain a fundamental understanding of how these decisions ought to be made, to develop methodologies and tools that can effectively support engineers in this task, and to largely automate the synthesis process. It is our conjecture that if we can make significant progress in addressing the above questions, our research work can greatly contribute to improving the quality of designs and reducing the time it takes to conceive them.

We are addressing three major questions in our projects. First, we are trying to understand how to deal effectively with the broad objectives and constraints in synthesis (e.g. economics, feasibility, downstream concerns). Second, we are trying to understand the role of qualitative and quantitative knowledge in the choice of design alternatives. Lastly, we are trying to understand how to deal with the very large and not well defined space of alternatives for synthesis. Ultimately, we are trying to develop and identify synthesis methodologies which are based on sound design principles, and therefore can cross discipline boundaries, and be applied effectively to major classes of design problems. As described in the next section, most of our projects have been aimed towards this goal.

The research work in the Synthesis Lab has concentrated in formalizing and understanding the following issues in each engineering domain:

1. **What are the appropriate representations for alternatives and objects to be synthesized?**

2. What are the basic approaches or design paradigms to tackle these problems?
3. What are effective search strategies for quickly identifying good or improved designs?

Accomplishments

Following is a brief summary of the major projects in the synthesis laboratory.

Preliminary Design of Engineering Systems In this project, Professor Maher (CE) has concentrated on an important synthesis paradigm. Synthesis is considered as a tree-search through a design space, where combinations of design decisions are identified that satisfy a set of constraints that are largely qualitative in nature (engineering principles or heuristics). The objective of the research at this point is to generate systematically all the feasible alternatives for a design. The accomplishments of this project are the development of EDESYN, a shell for building engineering design expert systems, and the use of that shell to build several expert systems for engineering applications.

EDESYN provides a domain-independent representation of the synthesis process as described above, and is implemented in a frame-based representation language running in Lisp. The synthesis process is represented as a constraint-directed search for feasible combinations of decisions incorporating planning to determine the relevant decisions. The knowledge base is defined by goals, or decisions, organized into levels of abstractions, mechanisms for satisfying goals, constraints on feasible solutions, and planning rules that determine the goals relevant to the current context. The design of a multi-window, graphics-based user interface has also been completed, and an evaluation method for selecting among the feasible alternatives is currently under development.

The shell EDESYN has been used to develop three expert systems: STRYPES, a designer of structural configurations for buildings; STANLAY, a designer for different layouts of a given structural configuration; and FOOTER a designer of foundations. STRYPES generates alternative combinations of materials and structural system types that are feasible for a given architectural plan. STANLAY, generates feasible layouts, and approximates the sizes of the structural components. FOOTER generates feasible foundation types and approximates the dimensions for given building loads and soil conditions.

The significance of this project lies in the ability to define quickly and test a representation for a design space for preliminary design. The three expert systems developed in the civil engineering domain were developed each over several weeks; this task took more than six months before EDESYN was available.

Expert Systems for Separation System Synthesis In this project, directed by Professor Westerberg, (ChE) expert systems are being developed for synthesizing a variety of separation systems. The main objective is to combine qualitative and quantitative knowledge, in order to quickly identify near optimal solutions. An opportunistic control structure is used for the search, which is based on a data-driven blackboard architecture.

Professor Lien of the Norwegian Technical Institute has worked jointly with Westerberg to develop AKORN-D, an expert system which synthesizes heat integrated distillation based separation systems and which has a completely data-driven control structure. This expert system has been tested successfully on several problems. For instance, in a five component mixture the search was reduced quickly to only two nearly equivalent solutions out of 1218 alternative designs.

Work has also been directed to the synthesis of separation systems with multiple mixed-products. An algorithmic method was developed which uses linear models for the columns, and where a bounding scheme is used to overcome the nonconvexities caused by the splitters. This scheme is able to identify the global optimal solution.

Another aspect of this project is the classification of separation problems for which distillation is not the likely separation technology. Work has been recently initiated on pressure swing adsorption. Work is also currently under way to develop a new expert system based on Knowledge Craft, and where the use of DPSK (Distributed Problem Solving Kernel) will be explored so the system can operate on several computers simultaneously.

This work is laying a foundation for integrating knowledge-based systems with analysis computations that are numerically intensive. Also, it is showing that this scheme allows the rapid identification of improved designs from a large number of alternatives.

Electro-Mechanical Design Environments Professor Rinderle (ME) has been investigating the design of various mechanical and electromechanical devices as a vehicle for evaluating various theses regarding the form-function structure of products. In particular, this work has evaluated the extent to which high-level form-function characteristics of components can guide design decisions at the conceptual stage, how these relationships can be identified and abstracted, and a synthesis strategy based on opportunistic combination of components characterized by their form-function structure. Also, the identification of critical design relationships has been investigated.

To identify form-function relations, devices are represented through a constraint network by equations and inequality constraints. These are then optimized parametrically in terms of a parameter that defines form (e.g diameter) through nonlinear programming. These computations yield interaction curves (e.g. torque vs. diameter in an electric motor), which due to changes in the active set of constraints, often exhibit non-monotonic behavior.

An environment for the conceptual design of mechanical systems is also under development and is being implemented in the program MEDA. MEDA is written in C and is an extension of M/P/E, a commercial CAD package. The basic idea is to provide a software environment that facilitates the successive aggregation of component models for electro-mechanical devices and automatically formulates the dynamic equations for motion.

Finally, the automatic identification of critical design relationships is also being studied. Through the use of transformations and ordering of equations, the objective is to not only facilitate their numerical solution, but also to provide better insight into the physical relationship of variables. A prototype program, EUDOXUS, has been implemented in Lisp to demonstrate the use of techniques for identifying critical design relationships.

The significance of this work is that it has started to provide a systematic framework for establishing function-form relations in electromechanical designs, as well as a framework for performing conceptual design in these systems.

Layout Synthesis This is a recent project by Professor Hemming (AR), who is investigating the systematic generation of alternative design topologies for two-dimensional layouts. Numerical optimization and constraint satisfaction techniques are being used to compute the geometry from a given topology. Applications are envisaged beyond the architectural domain.

This project is an outgrowth of the LOOS project that was initiated within the IBDE project of the Design Systems Lab. In LOOS two basic types of design variables are handled: spatial relations between rectangles, and upper and lower bounds on the corner coordinates of rectangles which are constrained by the topology. The former represent qualitative variables, and the latter quantitative variables. Generation and propagation rules are applied to develop configurations which are evaluated by a tester that combines linear programming and qualitative knowledge. Configurations are then evolved through a branch and bound procedure. LOOS has been applied to the design of configurations of kitchens and toilets.

The objective of the present project is to develop a general purpose layout synthesizer which could be applied for instance to VLSI design. The work will concentrate in the use of constraint-directed generation to avoid the enumeration of infeasible or unattractive solutions. Also a declarative language will be developed for the effective representation of geometrical constraints.

Automated Learning This project is being conducted by Professors Fenves (CE), Newell (CS) and Westerberg (ChE). The objective is to explore the potential of the SOAR system developed by Newell. This system has a number of intriguing features that may lead to a new generation of AI tools. SOAR provides two facilities not found in conventional expert system environments: a rich set of general or "weak" problem-solving methods built into the architecture; and the ability to "learn" by generating new chunks of knowledge from successful solutions of subproblems. Another interesting feature of SOAR is the fact that it is not committed to using a single search strategy which often leads to inflexible behavior.

One study in civil engineering implemented a small but representative problem for synthesizing feasible solutions satisfying input specifications and domain-specific constraints among the possible alternatives. The application dealt with selection of attributes of a floor slab for a given specification from a set of discrete alternatives with constraints. Experiments were run with different representations, and the resulting chunks and performance measures examined. The system exhibited learning within a task and across tasks. The study identified a number of strategies and representations to improve the quality of the learned chunks. On the basis of this exploratory study, a general research study is proposed with the aim of building lattice-like classifications or taxonomies of design specifications, descriptions and processes. The application domain is the design of bridges.

In the second study, synthesis of complex separation systems is being explored within the SOAR system. A major objective will be to develop an understanding of how the learning capability in SOAR handles engineering design problems where significant engineering computations (likely done in FORTRAN) must occur to analyze the performance of each system which is proposed. A general design strategy has been outlined within SOAR that will resort to numerical computation or optimization only when required. To accomplish this objective it is intended that the system will recognize "similar" problems and "good" designs. Furthermore, the solution will not be restricted by any simplifying assumptions (e.g. only distillation), and no commitment to an analysis strategy will be necessary. The first synthesis application that is being considered is the sharp separation of ideal mixtures without heat integration and will not involve engineering computations.

The significance of this project is that it may provide a new approach to synthesis and design with novel knowledge-based systems where intelligent use of qualitative and quantitative knowledge is made.

AI Applications

Synthesis is essentially a formative problem: the generation of alternatives subject to a set of constraints. Successful synthesis depends crucially on a number of AI techniques:

1. Representation of constraints, specifications and alternatives;
2. Abstraction of high-level design descriptions before the detail of the design are supplied;
3. Partial descriptions and evaluations of design concepts;
4. Declarative representation of qualitative information about partial designs and alternatives;
5. Search among feasible solutions; and
6. Planning of the search process.

The projects described each use several of the above AI techniques. Furthermore, without exception, the projects involve a combination of qualitative (symbolic) and quantitative (numeric) representations and evaluations. This is, without question, the unifying theme of engineering design applications of AI. It is not possible to address a realistic engineering design or synthesis problem using only "shallow" qualitative, symbolic representations of objects and knowledge. Rather, it is necessary to perform quantitative, numeric evaluations based on the underlying "deep" functional or causal relations.

Of the specific projects, the Separation Synthesis project is implemented in a "conventional" KBES framework (OPS5 and knowledge Craft). The Electro-Mechanical Design and Layout Synthesis projects aim at the declarative representation of the salient relations (form-function in the former, spatial in the latter) for reasoning and synthesis. EDESYN is our first attempt at a domain-independent synthesis "shell" — the design equivalent of EMYCIN. Lastly, our exploratory studies with SOAR are a first step toward a new generation of design expert systems with learning capabilities.

LIFE CYCLE ISSUES

Approach

We are developing strategies and tools to help the designer understand potential problems that might occur during the prototyping and manufacturing cycle of a product. Today, early design decisions frequently restrict the design space such that only suboptimal solutions regarding product manufacture are available. In many instances designers do not sufficiently consider appropriate design alternatives, since they are not aware that certain decisions may pose downstream problems. Inevitably, this leads to iterations, delays, and cost overruns.

Product designers need to worry about numerous candidate manufacturing processes simultaneously. They must consider quality, cost, lead time, availability of subcomponents, and much more. They must look at a product from a number of different perspectives. Most likely, designers make many decisions based on intuition because of time constraints or the unavailability of people who might have the necessary information. In addition, design alternatives can frequently not be evaluated because of time constraints.

The development of "simultaneous engineering" systems requires the understanding of two mutually orthogonal problems. One is concerned with fundamental representation issues, the other with the reasoning of domain specific process and product knowledge. Designing products and processes requires that one consider the functional specifications of an initial concept as well as the ultimate translation of these into actual 3D geometric entities. The human thinking process usually puts several auxiliary layers of abstraction between the functional specs and the 3D geometry. These serve primarily to facilitate the mental transition process in between. This process is assumed to be the same, independent of the level of detail at which the product or process is viewed. The underlying representation is expected to be valid independent of the resolution or application domain.

In product design, manufacturability is one of several major issues which may involve reasoning about forming, assembly, testing, cost, etc. If we pick forming as an area of concern, again one may have to reason about the specifics of processes like molding, cutting, stamping, and several others. While the product and process representation at least at the two lowest levels of abstraction (geometry and topology) can be shown to be the same, the manner in which specific information is processed is guaranteed to be largely different from one manufacturing process to another. Some processes can be described through mathematical models, while others may be too complicated for that. As an example, consider welding and molding. Welding turns out to be almost impossible to describe with tractable mathematical models; however, very important aspects of injection molding of plastic parts can be described with FEM models. The current focus of the representation development is geometry and topology. In the future we propose to take this representation to higher levels of abstraction to be able ultimately to perform the transition between functional specs and geometry. A new project addressing this problem has been recently initiated (See below).

Accomplishments

Our major accomplishments may be classified in two main categories: A unified geometric representation, and domain specific applications.

Unified Geometric Representation for Manufacturability Concerns

We first indicate the principles behind the underlying geometric and topological representation developed in the laboratory and then briefly discuss the manufacturability applications which are currently based on it. The key developer of the representation called 'NOODLES* is E. Gursoz (ME) who has received vital input from Professor Woodbury (AR).

Conventional approaches developed for solid modeling are restricted to solids only. In most instances, however, it is desirable to represent nonhomogeneous geometric entities (objects with different dimensionality) within the same data structure. The need for this capability arises from various reasons including yet incomplete models, modeling abstractions, and non-regular solids. Among existing solid modeling schemes, surface boundary representations are very promising for describing nonhomogeneous objects since the surface boundaries explicitly employ elements of lower dimensions.

Surface boundary representations in solid modeling traditionally have been based on two-manifold topologies (no more than two faces may coincide at one edge) [1]. Data structures and data manipulation formalisms introduced for two-manifold topologies cannot operate with non-manifold situations. Within the framework of non-manifold surface boundary representations, we have implemented a topology-rich geometric modeler which uses the node or corner, rather than the edge [2], as the fundamental element [3]. Furthermore, it was possible to realize non-regular set operators for nonhomogeneous entities. For example, intersections of two solids may render points, curves, surfaces as well as solids. Thus our non-manifold representation introduces new operations for model building in addition to conventional boolean operators. For example, solids may emerge as a result of adding new surfaces. NOODLES' very rich information structure regarding the adjacency and the use of all fundamental geometric elements lends itself naturally as a scheme to design with features and also for feature recognition.

Aside from forming the underlying representation of geometric objects for a number of EDRC projects, an immediate application of this work has been the prototyping of plastic parts at General Motors. The plastic components (lamps) are designed through highly developed surface modelers which can now be turned into solids through NOODLES. Once the solid model is built it can be passed to a stereo-lithography apparatus in which plastic parts are made through a laser-induced polymerization process.

NOODLES has been adopted as the geometric modeling tool for CMITs contribution to DICE (DARPA Initiative on Concurrent Engineering). Ultimately, it may form the representation for the entire DICE project. Furthermore, the constraint propagation network developed by Professor Rinderle (ME) will be integrated into the NOODLES scheme. We are working closely with West Virginia University (WVU) on the issue of supplying and supporting NOODLES. WVU acts as Coordinator among all participants of DICE.

Domain Specific Applications

Design for Molding The Design for Molding project is built around the NOODLES representation scheme. A set of form features which are considered important for molding such as bosses, ribs [4], and parting line are recognized through a graph-matching scheme. The features are geometrically decomposed. This frequently leads to the occurrence of incomplete 3D models. The decomposed features are evaluated regarding certain characteristic dimensions which then can be compared to proven manufacturing standards. The feature recognition scheme is furthermore useful for the automatic mesh generation which is required as input to commercially available flow analysis software tools (Moldflow, C-Flow, CAD Mold). All information regarding moldability is obtained through General Motors' Fisher Guide Division. An initial prototype of the Design for Molding software was demonstrated a year ago, and we are currently developing a plant prototype which GM intends to install in February or March of 1989.

NOODLES also forms the underlying data structure for a related effort on the calculation of 2D and 3D medial (symmetric) axis transforms MAT [5]. The motivation for this project originated from the need to simplify 3D geometric objects for topology extraction and feature recognition. We were able to demonstrate that recognition tasks may be significantly simplified after the application of a medial axis transform to a 3D object. The availability of this scheme is expected to facilitate the integration of manufacturing reasoning into Design for Manufacturability systems. The graph matching scheme described above (Design for Molding) for feature extraction will greatly benefit from this project. Currently, we have a stable implementation of MATs for arbitrary 2D objects and a less stable one for 3D ones.

Design for Stamping The Design for Stamping project is similar in structure to the Design for Molding project. The majority of the representation tools is used in both projects, whereas the domain specific reasoning obviously is not. Professor Desa (ME) is working closely with local industries to integrate relevant information. The current application focus is the stamping of I-C lead frames.

Design for Assembly The Design for Assembly project deals with the evaluation of a design for ease of assembly. Professors Khosla (ECE) and Sturges (ME) work together on this project. The objective is to create aids **for the** designer that will automatically evaluate the assemblability of a design. Given the model of the Mechanical System/Assembly (MSA) and **a** model of a set of assembly facilities, the proposed research will answer **the** question: Can the MSA be assembled automatically with **the** given facilities? Our approach consists of determining automatically a set of assembly operations, through a disassembly procedure, that leads to the given assembly (MSA). We propose to develop the theory and associated software tools that will reason about the assembly operations and address the question of assemblability. Again, the geometric modeling system NOODLES is being used to model the MSA. Abstractions of the MSA are in the form of the Component Graph (which explicitly provides the mating conditions between components), and the Component Hierarchy (which represents subassemblies in the MSA). Both of the above abstractions are created automatically. The MSA is divided into subassemblies based on the assembly procedure of these components. In most cases it has been seen that the assembly operations can be obtained by reversing the corresponding disassembly operations. Any assembly task can be considered to be an operation on a given set of components that constrains some or all of their degrees of freedom.

Single-Board Computer Configuration MICON is an expert system which selects and configures the components for a single board computer (SBC) based on user-provided functional specifications [6]. MICON has successfully configured working prototypes of several SBC's. Current limitations originate from the fact that MICON does not take into account the actual 3D geometry of the components in question. We have started to use NOODLES as the representation to check automatically for interferences as well as creating a FEM mesh for a heat transfer analysis. The key in future computer design systems will be the capability to perform electrical and mechanical analysis simultaneously. Professor Siewiorek (ECE) and E. Gursoz (ME) are cooperating to demonstrate the feasibility of such a design environment.

Manufacturing Cost Models Cost analysis is fundamental to any design project, since a designer, by making good choices, can decrease significantly manufacturing costs. We are attempting to develop a system which will give the designer realistic feedback about cost resulting from any incremental design decision. We have selected the Design for Molding project as a nucleus to develop general purpose cost models for use at all important design stages. Professors Datar, Kekre, and Mukhopadhyay (GSIA) together with researchers from the Design for Molding project defined a set of cost drivers which cumulatively determines the manufacturing cost of injection molded parts.

From observations at the GM plant, it was found that management of the overhead component of manufacturing costs is becoming increasingly important with automation, shorter product life cycles, increased product differentiation and faster customer response. The overhead burden as a percentage of total manufacturing costs in many plants exceeds 50%, with material costs accounting for 40% and direct labor only 10%. However, despite shrinking labor costs, allocation of the burden at various work centers is done on the traditional basis of either labor hours or machine hours. This leads to cross-subsidization of products and erroneous product costing, since labor hours or machine hours are not the true "cost drivers" of burden. Burden is caused by drivers such as set ups, inspection, material handling, number of parts, etc.

The models developed in this project permit the linkage of design decisions to cost drivers, and hence enable the designer to capture realistically all the costs incurred during the manufacturing cycle of the plastic parts (lamps): tooling, production of parts, inspection, assembly, finishing, etc. Currently, these models are not yet integrated into the software we are currently developing. We intend to do this over the period of the next two years.

AI Applications

As stated earlier, the successful introduction of life-cycle concerns into the early stages of design hinges on the understanding of two mutually orthogonal problems: a compatible representation of all issues of concern; and reasoning with domain-specific process and product knowledge. These two problems are also key concerns in AI. Thus, as in the Synthesis Lab, design research in the Design for Manufacturing Lab is intimately based on AI techniques for representation and knowledge-based reasoning.

The unified geometric representation embodied in the NOODLES system provides a consistent representation of geometry and topology suitable for reasoning over multiple levels of spatial abstractions. This representation is a major step in geometrical reasoning, in that it does not require that all components be represented to the same level of geometric rigor (i.e., as 3D solids).

The domain-specific projects (Design for Molding, Stamping and Assembly) all rely on the NOODLES representation. In addition, each of them uses substantial domain-specific, declarative knowledge bases for reasoning about their respective manufacturability concerns. The MICON system is a full-fledged knowledge-based design system hierarchically configuring a single board computer over a number of abstraction levels. All these projects use quantitative evaluations in addition to heuristic ones. The Manufacturing Cost Models project, although not explicitly Al-based, is intended to produce a very high level abstraction of manufacturing processes in support of design synthesis and evaluation.

SYSTEMS INTEGRATION

Approach

We are interested in improving the productivity of design systems, making it possible to generate better designs in much less time. The underlying premises are: (1) a product or process can be only as good as its design; and (2) a design can be only as good as the system in which it is produced. The implications are that manufacturers must maintain competitive design systems if they are to produce competitive products, and researchers must identify and eliminate systemic defects if they are to improve productivity.

It is convenient to think of a design system as consisting of:

1. Active or decision-making-agents. This category includes human designers and programmed tools.
2. Passive agents. This category includes modeling frameworks, knowledge bases, prototyping facilities and test facilities.
3. Organizations. This category includes structures for integrating agents and policies for managing their activities.

While the need for high quality agents (good designers, good CAD tools, good test facilities, etc.) has always been clear, the effects of organizations are only beginning to be understood [13]. For complex objects that require large design systems (such as cars, microelectronic chips, computers and buildings), organizational issues appear to be as important in determining productivity as the quality of the agents used. Nevertheless, the effort put into developing organizations for computer tools in design systems has been much less than the effort put into the tools themselves. The result is that much of the design effort is centered on tools—acquiring them, inserting them into design systems, learning to use them, collecting their inputs and interpreting their outputs.

Our objectives are to streamline the acquisition and use of tools, thereby speeding up the design process, leaving humans with more time to innovate, and ultimately to improve the quality of the designs. Our approach is to focus on six thrust areas that cover key organizational issues and certain new types of tools, as described below.

New Types of Tools Conventional tools concentrate on only some design activities of which drafting, simulation, analysis, optimization and synthesis are notable examples. Other important activities, such as supervising tools, managing information, coordinating tasks, building new tools, and upgrading design systems, are left for humans to do with relatively little automatic assistance. These gaps in automation limit productivity and must be closed if significant increases in productivity are to be obtained.

Tool Acquisition It usually takes a great deal of effort to add a new tool to a design system. Often, months or even years pass before designers become familiar enough with a tool to use it effectively. The research issue here is to develop ways for the much quicker absorption of tools into design systems, recognizing that these tools may be written in different languages and styles, and may prefer to reside in different types of computers.

Command and Control This area covers the specification of chains of command and reporting requirements. It also covers where and how decisions are to be made. For instance, are the activities of tool-A to be scheduled by tool-B, or is tool-A to be allowed to be autonomous, and if so, what conditions will induce it to undertake a task? Among the important questions that a control strategy must address are:

1. How are goals to be met?
2. How is backtracking to be done?
3. How are mistakes to be corrected?

Collaboration By collaboration, we mean the exchange of raw and processed data. These exchanges can occur in preplanned or spontaneous ways and can cover local or remote effects. Complex problems usually involve uncertainties and their solution processes are impossible to preplan completely. Instead, allowances must be made for mid-course corrections, which in turn, require mechanisms for the spontaneous (opportunistic) exchange of data. "Remote effects" are another important consideration. By "remote effects" we mean the impact of one design task on another. (For instance, the effects that the shape of an object can have on the design of its manufacturing process.) To make design-for-manufacturability, design-for-maintainability, etc. possible, one must set in place collaborative mechanisms for identifying and reducing any unwanted remote effects.

Human-Computer Interfaces The ideal interface would be easy to learn, easy to use, long-lived, universal in applicability (serve for a wide variety, if not all, tools) and allow humans to observe, intervene and control the activities of tools at any desired level of detail. Existing interfaces tend to be far from this ideal. Usually they are difficult to learn and are tool-specific. Moreover, they provide few facilities for interactively monitoring and intervening in ongoing computations.

Record Keeping and **Design Reuse** Since design is often an incremental process that relies heavily on previous iterations and projects, the success of a design effort is often determined by how easily information from previous efforts can be obtained and reused. This information includes not **only the** final results of previous efforts, but also supporting information **such as** explanations of how and why they were obtained.

In summary, our work is directed at ways to reorganize the computerized portions of design systems in order to make them more productive by:

1. Identifying and developing new types of tools for areas left uncovered by conventional types of tools.
2. Systematizing and streamlining tool acquisition.
3. Removing human-computer interfaces from individual tools and making the interfaces more general and stable.
4. Adding powerful, automatic facilities to control the activities of tools and promote collaborations among them.
5. Adding powerful record keeping and design reuse facilities.

Accomplishments

Our projects cover a number of domains, including microelectronics, construction, architecture, automobiles, power systems, chemical plants and steel making. Brief descriptions of some of the projects follow.

ASCEND (Advanced System for Computations in Engineering Design) In this project, Professors Westerberg (ChE) and Woodbury (AR) are concerned with the unacceptably long times it often takes to develop complex quantitative models in engineering, especially in areas of technology where no production modeling tools exist. A blocking problem is the lack of a powerful modeling language. The goal of the ASCEND system is to allow one or more engineers to develop and solve an engineering model consisting of thousands of equations an order of magnitude faster in terms of the engineer's time than is currently possible.

To meet these objectives, a language has been developed that supports powerful information structuring, and strongly partitions declarative information from procedural information. ASCEND currently consists of a programming language, a compiler, a structure browser, and two solving engines that permit engineers to develop, solve and/or optimize complex models comprising hundreds to thousands of algebraic equations more quickly than is currently possible.

Current effort is to add the capability to solve mixed sets of ordinary differential and algebraic equations. The user interface is being created with the Design Department in the College of Fine Arts. Three studies just starting are to extend the concepts to geometric reasoning, structured information gathering in preliminary design, and tolerance optimization in mechanical systems.

The significance of the ASCEND project is that assembling, manipulating and solving mathematical models is fundamental to much of engineering design. However, until now these activities have remained unsupported by computer aids in many engineering domains.

Geometric Reasoning In this project, Professor Woodbury (AR) is considering the problems of geometric reasoning. In much of design, decisions about the geometry of an object are a major constituent of the design process. Such decisions, whether analytic or synthetic, require deep knowledge of methods for manipulating geometry and a strong grasp of possible geometric configurations. To a certain degree, humans seem to perform geometric manipulations well, perhaps accounting for the current preponderance of human-driven, interactive interfaces in the geometric part of computer-based design systems. Available computational tools for geometry are less capable; they excel at representations within narrow formalisms, but can perform little analytic reasoning and have virtually no generative or synthetic abilities.

The purpose of the Geometric Reasoning project is to develop the necessary body of theory for sophisticated symbolic computation on geometry and to build convincing prototypes, applicable across a broad range of design applications. To this end, the project is pursuing the following set of research issues:

1. The generation and understanding of designs as the result of invoking a set of rules that describe parts of a design. Though fundamental to applying search-based strategies for design, formalisms for design based on physical objects (rather than drawings of those objects) have not emerged to date. Rule capabilities on both representations of individual parts and assemblies of parts are required.

2. Modeling of classes of designs in which variation between instances is captured by differences between continuous variables. This concept, which has been called variational geometry, requires a terse and elegant language if it is to become highly used in industry.

The project has produced a prototype implementation of a knowledge-based system framework that includes variational geometry capabilities, a two-manifold geometric modeling system and a grammar system for the design of structural components in buildings.

The significance of the project is in the fact that geometric reasoning is fundamental to much of design. As one would expect, the project supports and complements several other projects in the Center. Some of these links are listed below:

1. Rule-based geometry is a counterpart to the non-manifold modeling project. While the non-manifold modeling project is focused on data structures and query algorithms for the representation of complex design objects, the geometric reasoning project focuses on the automatic creation of models. Future work will merge both issues.
2. In the variational geometry effort the project uses the ASCEND modeling language as its basis. ASCEND provides a powerful and elegant means to describe complex geometric relationship.
3. Ongoing work on building representation complements work in the IBDE project.
4. Work on representing geometric uncertainties is being used by one of the critics in the CASE project.

CASE (Computer-Aided Simultaneous Engineering) In this project, Professors Talukdar (ECE) and Reddy (RI) are working on the design of automobile parts. In particular, they are developing "critics" to aid in the process of "simultaneous engineering." TTie term "simultaneous engineering" is used to mean technologies for controlling remote effects. Before one can reduce adverse remote effects to tolerable levels, one must have ways to identify and assess them. Tools called critics are used to make these identifications and assessments. A critic is, in essence, a self-activated analysis tool. The self-activation feature distinguishes critics from the common variety of analysis tools.

Designers cannot be expected to be aware of all the effects of their decisions. Many of these effects occur in areas with which they may have little or no familiarity. Therefore, even if analysis tools are available for these remote effects, designers are unlikely to know when and how to use them. Instead, they need self-activating analysis tools (critics) that are: smart enough to understand the decisions they are making, and capable of quickly providing them with feedback on the effects these decisions will have on remote tasks. Speed is needed because the benefits of feedback are greatest while the design decisions are still fresh and easily changed; the benefits are much less once the decisions have been allowed to solidify.

The goals of this project are: to develop methods for converting conventional types of analysis tools into critics by adding the capabilities of self-activation, and to demonstrate these methods in systems for the design of automobile parts. Progress toward these goals includes three critics that are nearing completion.

The significance of the CASE project is in the templates it will produce for critics in other domains. Without critics the goals of design-for-manufacturability, design-for-maintainability, design-for-reliability, etc. will remain unattainable in most engineering domains.

Towards a Unified VLSI Design Environment In the area of tool acquisition, this project, guided by Professor Director (ECE), proposes a more modular, object-oriented view of tools in which all tools have a front-end which is responsible for representing the abilities and status of the tool to the designer and to other tools. The creation of such an object-oriented view depends on an understanding of what a tool is, how it is to be used, what its requirements are, and how its results are to be employed. The project is working on methods to develop front-ends with this understanding and to tightly couple these front-ends with tools. The resulting objects will be much easier to insert into design systems than the original tools.

The problem of controlling tools is also helped by the object-oriented approach. By viewing tools as strongly typed objects, the representation of tools can be made more modular, allowing for generalized control procedures to be developed. Also, these control procedures will be more flexible and less dependent on a particular suite of tools. A prototype system, called CADWELD, that embodies the ideas on object-oriented tools and control procedures, has been completed.

In the area of interfaces, the project is developing unified command/help facilities which will assist users with the use and debugging of resident tools. The project is also working on ways to allow designers to easily create graphical displays. These are being implemented in a package called Mx. A prototype of Mx has been completed and is being used by tool builders. The current version of Mx has established the general Mx framework and includes menu, list, form and text windows. The design and implementation of the graphical-display-window type and data-linking capability is in progress.

The significance of this project is that its ideas (object-oriented tools, control procedures that are independent of tool suites, and interfaces that are independent of tools) are applicable wherever large numbers of design tools are used. Already, preliminary versions of Mx have been used by other tool builders, while the notion of object-like tools from CADWELD has been borrowed by CASE (a project described earlier).

A Framework for Knowledge Source Based VLSI Synthesis In this project, Professor Thomas (ECE) envisions design as the process of finding a path between two points or design states in a multi-dimensional design space. The initial design state consists of a system's functional specification and some set of market-driven constraints. The end point consists of a physical implementation that performs the necessary function **and** satisfies the constraints. Because of the complexity of the systems being designed today, a large number of intermediate design states at different levels of abstraction are traversed before reaching the end point. Utilizing this concept of a design space, synthesis tools can be envisioned as arcs in a directed graph that move the design from one node or design state to another.

This project has two goals. The first is to specify and demonstrate a framework that will integrate upper level VLSI synthesis tools by building on the graph model of design and on the salient features of the blackboard paradigm (tools should be as independent as possible from the information on which they operate; tools should be easily inserted and removed; and tools must have well defined points of communication and should be centrally managed). The second goal is to examine and generalize the process used to specify the framework in order to form a methodology which may be used in the specification of frameworks for product development environments in other domains.

Besides its contributions to the areas of collaboration and VLSI synthesis, the significance of this project lies in the graph model of design processes it is developing and the better understanding of design this model makes possible.

Flexible Mechanisms for Collaboration In this project, Professor Talukdar (ECE) is dealing with collaboration, control, interfaces and record keeping.

Collaboration (meaning the exchange of data among tools) can be achieved by a variety of mechanisms that have been developed in the fields of computer science, management science and human organizational theory. Which of these mechanisms are well suited to engineering design? Answering this question is the first goal of the project. The second goal is to select/develop a subset of good mechanisms for the computerized portions of design systems and to implement these mechanisms in ways that make them easy to use by designers and system builders.

Flexibility is a key property of a good collaboration mechanism because complex design processes involve great amounts of uncertainty and, therefore, cannot be completely preplanned. Tools that are hard-wired together, or tools whose interconnections can only be reconfigured prior to computations, have trouble with processes that cannot be preplanned. If tools are to effectively participate in handling uncertainties (such as design changes and other midstream changes of direction) they must be provided with more flexible means for collaborating.

As the project finds good mechanisms for collaboration, it implements them in ways that are easy to use. The result has been two completed sets of collaboration aids and a third that is in development. The first, called **COPS,** helps system builders to construct multiple blackboards in a network of computers. Thereby, system builders can easily obtain the advantages of the blackboard paradigm, namely, tools that are independent of the data on which they operate, and increased ease with which tools can be inserted and removed from systems.

The second set, called DPSK, allows system builders to quickly construct object-oriented data structures that can be distributed over a network of computers and shared by an arbitrarily large number of programs. Thereby, DPSK provides convenient ways to integrate programs written in different languages and housed in different computers. DPSK provides not only the blackboard based collaboration mechanism of COPS, but also some mechanisms borrowed from human organizational theory. Manuals and source code for DPSK are available, and are expected to be used by an expanding cadre of projects both inside and outside the Lab.

The third set of aids, called FORS (Flexible Organizations), is under development. It uses a network-based model of collaboration that is more general and seems to cover the needs of design projects better than the models underlying COPS and DPSK. The network model was arrived at by studying design processes in the automobile and power systems industries. Nevertheless, it bears a striking resemblance to models for VLSI design being considered by Professor Thomas (see previous description) and models for computers being considered by Professor Siewiorek (in the Design for Manufacturing Lab).

Besides implementing the network model of collaboration, FORS allows for the dynamic reconfiguration of tool interconnections, provides an icon-based interface that is virtually self-explanatory and maintains historical records and activity traces. A prototype of FORS with a small suite of tools for designing certain automobile parts has been completed. Work on adding tool suites from other domains is in progress.

The significance of this project is twofold. First, it is providing guidelines and generic aids for designers and system builders to use in connecting and reconnecting tools. Second, it is providing models by which we can better understand engineering design.

Information Structuring for Preliminary Design In this project, Professor Westerberg (ChE) is dealing with new types of tools and record keeping. The project began in July of 1987 when teams from Westinghouse Electric Corporation and EDRC initiated a project to study the design process. The project was for a team from the EDRC to observe the design of a new type of control system for coal-fired electric generation plants. Because the Westinghouse team was unfamiliar with die required methodology, it had to develop an information base on which to develop the design concept. Based on its observations, the CMU team prepared a white paper to describe potential new aids which could improve the activity. One of these was a tool permitting information structuring/modeling to support the early stages of information gathering. This tool formed the basis for the second year of activity which started in July, 1988.

Humans will be the agents to search the "world" knowledge base available, to locate and assess the information relevant to the design. The proposed tool is to aid in the sharing and use of that information with others. Initially the information is entered into the computer in the form of memos or as copies of overhead transparencies, i.e., in the form that information is currently being exchanged. This information is kept in a hypercard-like database which can be structured and searched as desired. "Refinements" of each memo can be tagged by any member of the design team to identify key concepts, variables and relationships contained within it. Consistency is not checked here.

We argue that a design team works to construct a model of the design, starting at a high level of abstraction and working slowly to a more detailed representation. We are providing a language in which this modeling can be done, with constructs that permit it to be tied into the memos and their informal refinements. This more formal model can be checked for correctness in powerful ways. At the highest level, it will be similar to a semantic network of concepts. At the lower levels it will contain quantitative relationships that can form the basis for computations on the model. The language is being motivated by our work with ASCEND and by knowledge representation schemes used in AI. With its ties into the informal information, the formal model operates as a key for searching that information.

Westinghouse employs a methodology for structuring models that is similar to the above in intent, except that it has not been formalized. We are sharing ideas with them and working together to create this new modeling capability.

The significance of this project is that it will provide a set of concepts and a prototype tool for supporting the early stages of a "new" design project.

EBDE (Integrated Building Design Environment) In this project, Professors Fenves (CE), Flemming (AR), Hendrickson (CE), and Maher (CE) are focusing on the integration of activities of the various professional groups engaged in the design and construction-planning of a building; and providing a testbed for the exploration of ideas from other **projects.**

Seven interacting knowledge-based processes **are** involved in **the IBDE** project They are:

1. ARCHPLAN, an architectural planning expert system which assists in the conceptual design of a building;
2. CORE, which generates layouts of the elements in the service core of the building under consideration (elevators, elevator lobbies, restrooms, emergency stairs, utility rooms, etc.);
3. STRYPES, which configures a structural system for the building;
4. STANLAY, which develops a layout of the structural system specified by STRYPES and then performs an approximate analysis of the lateral and gravity load-resisting systems to determine the load requirements for component design;
5. SPEX, which performs the preliminary design of the structural components for the structural system specified by STANLAY;
6. FOOTER, which makes a preliminary design of the foundation of the building; and
7. CONSTRUCTION PLANEX, which assists the construction planner.

A prototype that vertically integrates the seven processes has been successfully demonstrated. Issues of feedback and conflict resolution are being explored. The current communication and control mechanism is implemented using DPSK. The next version may be built on top of **FORS.** The possibility of using or adapting the interface facilities from other lab projects is being explored. The geometric reasoning and synthesis functions of the processes are closely coupled to related EDRC projects.

The significance of the IBDE project is twofold. First, it is has the potential for considerably improving productivity in the design of buildings. Second, it is the only system within the Design Systems **Lab** that is large and general enough to serve as a testbed for ideas and software generated by other projects.

AI Applications

The interaction between Design Systems Lab activities and AI techniques is best illustrated in terms of four clusters of primary emphasis.

The ASCEND and Geometric reasoning project concentrate primarily on representation issues. ASCEND deals with the representation of complex quantitative relations in engineering models; it relies heavily on declarative representation and on the exploitation of powerful relations among entities. The geometric reasoning project deals with representation of geometry in support of spatial reasoning in rule-based and grammar-based design systems.

The VLSI synthesis framework and the Building Design Environment deal primarily with cooperative problem-solving involving multiple knowledge-based systems communicating through a blackboard architecture. They both address issues of design evolution, backtracking and constraint propagation.

The Unified VLSI design environment and the Flexible Collaboration Methods project are primarily concerned with the problems of design tool abstraction and the control of such tools. In both projects, the major objective is to represent the salient characteristics of design tools, interfaces and design data so that a wide range of control strategies (manual and automatic) may be deployed.

Finally, the CASE and Information Structuring projects are primarily addressing process abstraction: the representation of the design processes at high levels of abstraction so that decisions made by one design agent in the design system can be automatically propagated to the other agents, and remote effects, constraints, etc. resulting from the decision are "fed back" to the original agents.

SUMMARY

We are learning much about design. Most recently, we have attempted to integrate our experience to date into a classification system for design problems that includes the dimensions of:

1. original design *versus* routine design;
2. amount of activity required (large: Boeing 747, *versus* small: doorknob);
3. strong *versus* weak coupling in the decompositions used;
4. degree of importance of geometry;
5. quantity produced; and
6. domain of the design.

Boundary conditions, such as the availability of outsourcing parts for

the product, labor costs, and so forth, are also a part of this classification system. It is an important component of our planning, as it shows where we need new activities. We see a cycle of deductive and inductive reasoning activities in our planning, with a time constant of approximately two to three years.

In its motivation, staffing and methodology, EDRC is an *engineering* research center, and not a "pure" AI applications center. However, as the presentations indicate, there is an intimate match between the concerns of EDRC and the concepts, methods and research interests of AI. Therefore, EDRC has made a determined effort to understand, import, apply and expand AI methodologies to serve its purposes. Our objective is not to be at the cutting edge of AI research but, rather, to evaluate and apply cutting-edge AI research to engineering design needs (as well as cutting-edge research in other computer-based methodologies, including optimization, databases, programming languages and systems, distributed processing, etc.). We believe that we have been successful in this objective: AI concepts and methodologies pervade a significant portion our projects, and several of our projects, such as the SOAR and CASE projects, are intimately linked to the most advanced AI developments.

Our success of integrating and "internalizing" AI is best exemplified by the resignation letter of our first research scientist, Michael Rychener. In that letter, Mike points out that initially he could play a central role in educating students and staff about AI methodologies and assist in the definition of several projects by identifying appropriate AI methodologies and paradigms. However, Mike felt that as the understanding and use of AI methodologies began to pervade EDRC, the need for his expertise decreased to the point where it ceased to be useful This anecdote exemplifies our success in adopting AI methodologies to our objectives.

REFERENCES

1. **Weiler, KJ.(1985). Edge-Based Data Structures for Solid Modeling in Curved-Surface Environments, IEEE CG&A, January.**

2. **Weiler, KJ. (1986). Topological Structures for Geometric Modeling, Ph.D. Thesis, Rensselaer Polytechnic Institute.**

3. **Gursoz, L.E. and Prinz, F.B. (1988). Corner-Based Representations of Non-manifold Surface Boundaries in Geometric Modeling, Proceedings of the RPMFIP-NSF Workshop on Geometric Modeling, Rensselaerville, NY.**

4. **Nackman, D.R. and Pizer, S.M. (1985). Three-Dimensional Shape Description Using the Symmetric Axis Transform I: Theory, IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. PAMI-7,No.2: 187-202.**

5- Dixon, J.R. (1988). Designing with Features: Building Manufacturing Knowledge into more intelligent CAD Systems, Manufacturing International 88, Atlanta, GA.

6. Birmingham, W.P., Brennan, A., Gupta, A.P. and Siewiorek, D.P. (1988). MICON: a single Board Computer Synthesis Tool, IEEE Circuits and Devices Magazine.

7. Westerberg, A.W. (1988). Synthesis in Engineering Systems, EDRC Report 06-47-88, Carnegie Mellon University,

8. Fraser, D.J. (1981) Conceptual Design and Preliminary Analysis in Structures, Pitman Publishing, 1981.

9. Nishida, N., Stephanopoulos, G., and Westerberg, A.W. (1981). Review of Process Synthesis, AlChE Journal 27:321.

10. Douglas, J.M. (1988). Conceptual Design of Chemical Processes, McGraw Hill.

11. DeMicheli, G. and Sangiovanni-Vincentelli, A. (eds) (1987). Logic Synthesis and Silicon Compilation for VLSI Design, NATO Advanced Study Institute, ASI series-E, Volume 136.

12. Finger, S. and Dixon, J. (1989). Mechanical Engineering Design Theory and Methodology: A Research Review, to appear in Research in Engineering Design, 1:1.

13. Clark, K.B. Chew, W.B., and Fujimoto, T. (1987). Product Development in the World Auto Industry: Strategy, Organization, and Performance, Harvard Business School.