REPRINT

Metaphors for Manufacturing: What Could it be Like to be a Manufacturing System?

John L. Casti

THE INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS

is a nongovernmental research institution, bringing together scientists from around the world to work on problems of common concern. Situated in Laxenburg, Austria, IIASA was founded in October 1972 by the academies of science and equivalent organizations of twelve countries. Its founders gave IIASA a unique position outside national, disciplinary, and institutional boundaries so that it might take the broadest possible view in pursuing its objectives:

To promote international cooperation in solving problems arising from social, economic, technological, and environmental change

To create a network of institutions in the national member organization countries and elsewhere for joint scientific research

To develop and formalize systems analysis and the sciences contributing to it, and promote the use of analytical techniques needed to evaluate and address complex problems

To inform policy advisors and decision makers about the potential application of the Institute's work to such problems

The Institute now has national member organizations in the following countries:

Austria
The Austrian Academy of Sciences

Bulgaria
The National Committee for Applied Systems Analysis and Management

Canada
The Canadian Committee for IIASA

Czechoslovakia
The Committee for IIASA of the Czechoslovak Socialist Republic

Finland
The Finnish Committee for IIASA

France
The French Association for the Development of Systems Analysis

German Democratic Republic
The Academy of Sciences of the German Democratic Republic

Federal Republic of Germany
Association for the Advancement of IIASA

Hungary
The Hungarian Committee for Applied Systems Analysis

Italy
The National Research Council

Japan
The Japan Committee for IIASA

Netherlands
The Foundation IIASA–Netherlands

Poland
The Polish Academy of Sciences

Sweden
The Swedish Council for Planning and Coordination of Research

Union of Soviet Socialist Republics
The Academy of Sciences of the Union of Soviet Socialist Republics

United States of America
The American Academy of Arts and Sciences

METAPHORS FOR MANUFACTURING: WHAT COULD IT BE LIKE TO BE A MANUFACTURING SYSTEM?

John L. Casti
Theory of Manufacturing Project, International Institute for Applied Systems Analysis, Laxenburg, Austria

RR-86-4
October 1986

Reprinted from *Technological Forecasting and Social Change*, 29, 241–270 (1986).

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
Laxenburg, Austria

Cø Bø

Reprinted with permission from *Technological Forecasting and Social Change*, 29, 241–270 (1986).

Printed by Novographic, Vienna, Austria

FOREWORD

During the course of the Theory of Manufacturing Feasibility Study, one of the topics of interest was to investigate the likelihood of a system-based framework for the development of a general theory of manufacturing processes.

In this report, John Casti presents several such candidate structures, using the argument that it is best to approach the question of theoretical structures for manufacturing by means of "metaphors". Here the metaphors are taken from engineering, biology, linguistics, computer science, and chemistry, thereby providing a menu of alternative starting points for a research program addressed to the creation of a theoretical foundation for manufacturing.

THOMAS H. LEE
Director
International Institute for Applied Systems Analysis

TECHNOLOGICAL FORECASTING AND SOCIAL CHANGE 29, 241–270 (1986)

Metaphors for Manufacturing: What Could it be Like to be a Manufacturing System?

JOHN L. CASTI

ABSTRACT

In the world of modern manufacturing, considerable emphasis is placed upon properties of manufacturing systems described by the terms "flexibility," "complexity," "reliability," "self-repairing," and so forth. To understand and deal with such properties, one needs a theoretical framework allowing one to pose and analyze various questions surrounding the meaning and interrelations of these terms. This paper addresses the questions of what sort of system-theoretic frameworks would be likely candidates to form the basis for such a *theory of manufacturing*.

Our approach to the modeling problem is to examine several different paradigms that have proven useful in other fields—engineering, biology, linguistics, computer science, chemistry—and to explore the degree to which these "metaphors" can be used to characterize manufacturing systems. The paper concludes with an assessment of the strengths and weaknesses of each metaphor and a suggestion for a program devoted to further development of the most promising approaches.

I. Manufacturing Systems—What's All the Fuss?

The great American philosopher, and sometime baseball player, Satchel Paige, once remarked when asked his secret for a long life, "Don't ever look back, 'cause somethin' might be gaining on you." The 1950s and 1960s saw U.S. and European manufacturers following Satch's dictum religiously, perhaps in the hope that long life for baseball players and long life for manufacturing enterprises were somehow subject to the same natural laws. This was not to be. The widespread availability of cheap, special-purpose computing machinery coupled with an aggressive and innovative management and production attitude in Japan (and elsewhere in Southeast Asia) caught the complacent Western manufacturers looking the other way, resulting in the current explosive corporate soul searching centering on the theme, "Are we losing our technological edge?"

Naturally, the sociopsychological corporate insecurities would have emerged without the added impetus of the challenge from Japan. The main force powering the wholesale re-examination of manufacturing was the arrival of unprecedented computing capability, not the Japanese. They were only the first to exploit this capability in a widespread manner. All of the traditional procedures for transforming one set of physical objects into another, starting with design and going on to final distribution and consumption, have

JOHN L. CASTI is a Research Scholar at the International Institute for Applied Systems Analysis, Laxenburg, Austria.

Address reprint requests to Dr. John L. Casti, IIASA, A-2361 Laxenburg, Austria.

 0040-1625/86/$03.50

been called into question by our new-found ability to cheaply, quickly, and accurately process information. This is the essence of the identity crisis currently under way in the world's manufacturing industries: a need to completely rethink the entire manufacturing process from top to bottom, with a view toward putting it back together in a new structure suitable for taking advantage of enhanced information processing and automation capabilities. But what are the characterizing features of this newly emerging structure (or structures)?

The most distinguishing aspect of the manufacturing plant of the future is its *heterarchic* structure. The ability to transfer information almost instantaneously from one part of the process to another means that the traditional hierarchical, tree-structured, sequential manufacturing system will be transformed into a heterarchical, distributed, parallel-processing system capable of a high level of flexibility in producing a myriad of products with high quality and efficiency. The effective coordination of such a distributed process would be unthinkable, of course, without the information-processing resources that have recently become available. It is probably not an exaggeration to say that the ultimate aim of any decent theoretic study of such manufacturing processes is to devise a framework that enables us to understand how to reconfigure the various components of the manufacturing process (design, production, distribution, management) to most efficiently and effectively employ the computing resources available. Let us consider a typical scenario for the factory of the future.

The process of designing a product will be an iterative interaction between the designer and the computer (CAD). The designer will supply the product, concepts, and specifications, and the computer will carry out design calculations and provide standardized information. During this process, the computer can be continually taking into account information on the manufacturing costs and capabilities needed to produce the product under design. The computer then employs this information to generate a design that not only meets the product specifications, but also can be manufactured in some "optimal" way. This design phase of the process may be physically far removed from the actual plant facilities involved in the production of the product. Nevertheless, information technology now enables the design computer to be continually apprised of the status of the plant and to employ this information as part of the design process.

At almost the same time as the design process is going on, the production-planning part of the system is using the design information to set up an optimized plan for production of the product. This plan involves selecting the proper equipment and processes, configuring the sequence of operations, choosing the operating conditions, and so forth. All of the design and production information is then used to control the automatic machines that actually do the manufacturing. Each of these machines continually feeds information about its status back to the production-control system, which then performs dynamic adjustments to the production plan as needed.

While the production process is under way, the various machines carry on self-diagnosis of their condition. If a failure is impending, they take automatic corrective actions. In addition, the machines also carry out automatic quality-control inspections at each stage of the product's manufacture, so that the final finished product is fully inspected and conforms to the original design requirements.

During the course of production, the distribution-planning component of the system is in communication with the production part, gathering information about how to best allocate the finished product among various distribution centers. The distribution program must optimally balance current demands and order backlogs with available transportation facilities and costs to decide the optimal means for distributing the finished product among various distributors, consumers and inventory warehouses.

This brief, skeletal outline of the operation of the factory of the future is notable for its reliance upon a high degree of communication, both within each major component and, more importantly, *between* components. Each stage must plan its action upon information about what happened in the preceding stages. The loop is closed by management strategists who will analyze the results of the distribution network (sales, profits, and the marketplace) and feed this information back to the design stage. The theoretic frameworks introduced later are all designed to explicitly accommodate this essential feedback-feedforward characteristic of modern manufacturing concerns.

II. Manufacturing As a Systems Problem

Before embarking upon a consideration of manufacturing processes from a system scientist's point of view, let us define the term manufacturing system. For us the following equation holds:

$$\text{manufacturing system} = \text{design} + \text{production} + \text{distribution} + \text{management}.$$

As will be seen, each level of the manufacturing process contains its own distinctive mixture of problems and concepts addressable by various system-theoretic ideas and tools. Sometimes the same system ideas can be employed at different levels; more often they cannot. One of the first steps toward the development of any theory of manufacturing is to isolate the level at which the problem is being entered and to identify the system framework and tools most suitable for addressing the characteristic problems of that level. We shall take an initial step in this direction in later sections, but for the moment let us look at a stratification of that multifaceted object termed a manufacturing system.

In Figure 1, seven distinct levels are identified at which it seems to make sense to consider the manufacturing system as a distinguishable entity. The lowest such level is that termed Parts, which is concerned with the manufacture of individual parts of some

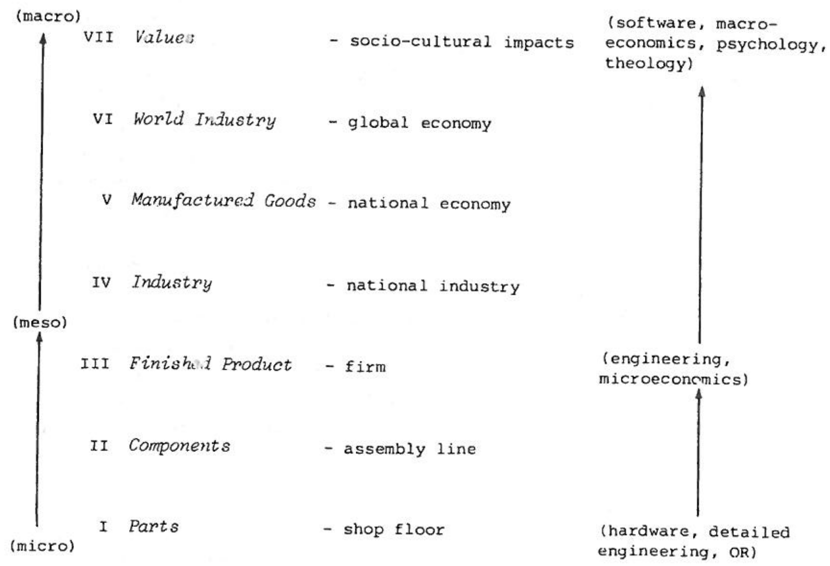

Fig. 1. Manufacturing system stratification.

finished product. At the other end of the spectrum of levels is that labeled Values, which involves itself with the sociocultural impacts of the manufacturing process.

Roughly speaking, each level is composed of a set of elements from the level below, with the exception of the first and last levels. Also, as one moves from Level I to Level VII, a rather smooth transition occurs from hardware and physical-science dominated issues to matters of software and social and human science. This stratification is by no means cast in concrete; many alternative breakdowns are possible. For example, one could expand Level II (Components) into sublevels such as materials, parts, and subassemblies. Because only the *idea* of a stratification is important here, and not the particular choice of levels, we shall refrain from exploring other possibilities in this paper.

Each level of stratification contains its own set of basic questions that make up the foundations for a research program for manufacturing systems, *at that level*. For instance, at Level I (Parts) the basic questions center on the age-old issues of scheduling, flow planning, and inventory control (protection stock). Typical issues are how to integrate the scheduling and planning activities (i.e., how to create a Gantt chart schedule that is consistent with the flow/production plan) or how to spread protection stock across the various levels of the bill of material. Roughly speaking, the basic questions at this level are the same as they have always been, suitably "souped up" to account for the new cross-activity communication possibilities offered by large-scale computerization.

On the other hand, a totally different set of questions arises when we look at manufacturing as a purely economic process as in, say, Level V (Manufactured Goods). Here one is concerned with the role that manufactured goods play in the overall economy of a nation, and the main questions revolve about how manufacturing as an economic activity interfaces with other components of the economy, such as the service sector and agriculture. In these instances, the effects of computerization in manufacturing systems are seen only to the extent that they directly impact the influence of manufacturing in the overall economy.

If one moves to an even higher level, such as Level VII (Values), then the emphasis on economic questions fades away and is replaced by central issues involving the influence of manufacturing systems upon the individual's sense of well-being, the way in which new manufacturing systems change the overall quality of life, and so forth.

Happily enough, a diligent systems scientist can very likely find grist for his mill at each level. Those with an analytic orientation can direct their attention to the operations research-oriented questions at the microlevels, say Levels I–IV, whereas philosophically inclined systems thinkers will probably be happiest devoting their speculations to the issues at the higher Levels VI–VII. Levels IV–V are probably of interest mainly for economists.

This article will examine the middle-of-the-road Levels II and III; an area somewhere between the nitty-gritty problems of very specific algorithms and techniques for very specific manufacturing systems and the almost ethereal problems of macroeconomics and manufacturing. Basically, the interest here is at the level of the *firm*, which roughly comprises the activities of design, production, distribution, and management planning. In order to keep the discussion within reasonable bounds, we shall concentrate only upon the production activity as a point of focus for the metaphoric models to be developed later.

III. Production Systems in the Computer Age

A theory of manufacturing, even at the restricted level of an individual firm's production system, already entails the capacity to effectively address a broad range of new systems issues arising from the computation and communication capabilities offered

by modern technology. Before outlining any modeling metaphors upon which to begin development of an axiomatic theory of manufacturing, let us examine some of the more important features of production that such a theory should encompass.

As already noted, the most distinguishing characteristic of a modern manufacturing process is its high degree of *decentralization;* i.e., its heterarchic structure. In the production of a given product, several components are produced simultaneously, often at different locations, and an important part of the theory of such production involves establishment of effective communication linkages and scheduling procedures to insure integration of the overall process. Determination of good communication-network structures and means to accommodate time lags in the communication of both information and materials are an essential aspect of any theory of production.

Another key system issue is *flexibility*. This is a vastly overworked and, as yet, ill-understood aspect of manufacturing processes. Everyone seems to agree that flexibility in production is a good thing and that if one doesn't have it one should try to get it, but a singular silence falls when researchers come to the point of identifying exactly what constitutes a "flexible" production system. An axiomatic theory of production should provide a means for comparing two production processes and for rendering some sort of verdict as to whether one is more flexible than the other. Related issues such as the interconnection between a production system's flexibility and its complexity and reliability are also integral to any decent theory of production systems.

Questions of *reliability, automatic fault detection,* and *automatic repair* will also play an important role in the production systems of the future. A theory of such processes must be able to incorporate various feedback-feedforward mechanisms that will enable each component of the production process to continually monitor its performance and to self-adjust its operation to ensure uniform quality of its output. From a theoretical point of view, this will entail data paths and components in the model of the production process that are devoted to the automatic maintenance and repair of the individual manufacturing cell. Such mechanisms are, of course, well known in biology, but at present we know little about how they can be included in large-scale manufacturing systems.

An aspect of system flexibility, one important enough to be singled out for individual attention, is the property of *self-organization*. The production system of the future must have the capability to automatically reconfigure itself in order to adjust to changing environmental circumstances. Disruptions in raw-material supplies, breakdown of a system component, rush orders, and so on are endemic to any production environment, and an automatic-production system must be able to immediately react to such perturbations by rearranging its components and communication paths to continue functioning despite these disturbances. This means that a certain level of redundancy is needed in the system to provide the means for such self-organization to take place. A good theory of production will provide a means for deciding what kind of redundancy is needed, how much is needed, and where it is needed to ensure the system's survivability in the face of such unexpected and unpredictable "shocks."

Finally, we come to the problem of *learning and evolution*. In many ways, an automatic-production process is like a living organism: it must "metabolize" inputs, produce outputs, and engage in continual self-monitoring and repair. This picture also suggests that such a production system must also learn about its environment and adjust its internal structure and workings to best fit the circumstances it finds itself in; that is, the system must learn, adapt, and evolve. If we accept this argument, then a theory of such processes must include a learning and adaptation mechanism that allows the production system to adjust itself to fluctuating external operating conditions. Again, like self-organization, such properties are another facet of the overall issue of flexibility but with the distinguishing feature that they involve *implicit* system properties of the production process that only make themselves explicit as the process unfolds and additional information becomes available.

In the attempt to provide modeling frameworks suitable for encompassing the foregoing system aspects of production processes, we shall explore a top-down, functional approach to the overall production system. By this we mean that the individual components of the production system will be identified by the *functions* they perform in the process and will not be characterized by any of their structural properties (size, weight, cost, color, and so on). So, for example, if welding is required in the process, one component may include an object termed "welder," with its only features being its input and output channels, together with its functional role in the overall process. Such an approach focuses upon the information content and flow through the production process and tends to deemphasize the details associated with specific machines, computers, and communication paths. We contend that the functional orientation is the most natural way to get at the systems aspects of production and that once the functional framework is understood it can then be translated into the necessary structural terms. This approach is much the same as what is done in biology when we disregard the individual biochemical reactions and speak of the cell's metabolic and genetic activities in functional terms, with the chemistry reintroduced after the functional set-up is made clear. The premise here is that the same idea can be used to investigate the theory of production in a manufacturing system.

IV. Natural Systems and Formal Models

Before embarking upon a presentation of a modeling framework suitable for constructing a theory of manufacturing processes, let us pause for a moment to consider the even more general issue of what constitutes a model of a given natural phenomenon and the means by which we encapsulate some subset of the real world within a mathematical representation.

The basic schema underlying all mathematical modeling efforts in any field is depicted in Figure 2.

Here we have some natural phenomenon N, a production system in our case, characterized by some collection of physical observables $f = \{f_1, f_2, \ldots, f_n\}$; for example, raw-materials inputs, machinery of different types, or computing resources. These observables are usually taken to be real-valued mappings defined upon some underlying abstract set of states Ω; that is,

$$f_i : \Omega \to R, \qquad i = 1, 2, \ldots, n,$$

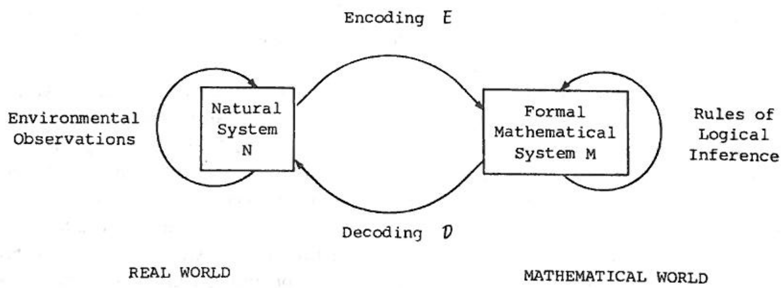

Fig. 2. General schemata for mathematical modeling.

although generalizations are possible. As a result of experimental observations and/or other information, certain mathematic relationships ϕ_i are found linking the values of the observables. We term such relationships *equations of state* and write them as

$$\phi_j = \phi_j(f_1, f_2, \ldots, f_n) = 0 \qquad j = 1, 2, \ldots, m,$$

where m is the number of such relations. Thus, on the left side of Figure 2, we have the data of our system N encapsulated as a map

$$\phi : R^n \to R^m.$$

Roughly speaking, this is an input/output, or phenomenologic, description of N. We would now like to formulate a mathematical model of this situation within some formal mathematical structure, in order to provide a means for making predictions about the behavior of N. It is important to emphasize the point that the phenomenologic description in terms of ϕ has no predictive or causal structure; it is just an account of experimental observations.

Turning to the right half of Figure 2, let us examine what is meant by a "formal" mathematical system. Speaking informally (all punning aside), a formal system comprises four components:

1. An *alphabet* of abstract symbols from which we compose statements of the system;
2. Rules of *grammar,* that express those symbol strings that are admissible statements;
3. Logical *rules of inference,* which allow us to generate new admissible symbol strings from old ones;
4. A set of *axioms,* which are given symbol strings assumed to be admissible without proof.

Note that a formal system has no intrinsic *semantic* content. The symbol strings are invested with no particular meaning or interpretation; they are merely collections of abstract objects. This point enables us to employ the same mathematical structure in widely differing areas by attaching different meanings to the symbols as the situation of interest dictates.

A good example of a formal system is a natural language like German or English, which makes use of compound words. The symbols of such a system are the letters of the alphabet, and axioms are the words of the language. The grammar of the language tells us the admissible patterns of stringing words together, whereas the rules of inference allow us to generate new words from old by, for instance, combining two words such as *Hausmeister* or babysitter.

A somewhat more fanciful example from the manufacturing area would be to regard the individual parts available for building a product as constituting the alphabet of symbols for the formal system. The symbol strings are then merely collections of parts or assemblies ordered, say, by the sequence in which the parts must be assembled. The axioms would correspond to assemblies already available; that is, something like standard component assemblies that the firm has on hand. The grammar would be a specification of the ways in which parts can be put together to form assemblies; that is, to a physically realizable scheme for combining parts. Finally, the rules of inference constitute a set of operations that can be performed upon assemblies to create new assemblies. Roughly speaking, such rules correspond to the rules of design. Carrying this example to the level of semantics, the interpretation of what the assemblies *mean* seems to correspond to the idea of the functional characteristics of the product; that is, the meaning of the assembly is what *function* the object is intended to perform.

Returning now to the modeling relationship shown in Figure 2, the encoding and decoding operations E and D, which lie at the heart of any mathematical modeling effort, play the role of mediators between the real world on the left side of the diagram and the mathematical world on the right. The observables of the real world and the equations of state are encoded by E into the objects of some formal system. The logical rules underlying the formal system are then employed to derive true mathematical statements (theorems) from the axioms of the system. Finally, the theorems are decoded by D and interpreted as statements about the natural system; these are the essential steps in establishing a modeling relationship between the natural system N and the formal mathematical system M. If all the elements (the observables, E, D, and M) are chosen adroitly, then the theorems of M when decoded will correspond to true statements about the observables of N; that is, the theorems will *predict* new relationships among the observables. For the purpose of uncovering these new relations we enter into a modeling exercise, and, as the foregoing discussion makes clear, the only reason for using a mathematical model is the natural predictive interpretations that can be attached to the process of formal, logical inference. Without this "automatic" means to generate new facts (theorems) from old, no particular advantage would be gained by going through the often difficult and distorting encoding and decoding operations E and D. In the following section, we shall attempt to justify the abstract schema of Figure 2 by presenting a number of plausible formal models that appear to hold promise for acting as a foundation upon which to base a theory of manufacturing for a firm (Level III in Figure 1).

V. Some Metaphors to Build By

Because at the moment we have no clear-cut view of what constitutes the "right" mathematical paradigm for representing the production process in a firm, the modeling frameworks presented below have all emerged as answers to the question, "What could it be like to be a production process?" Thus, each of the paradigms sketched here is a *metaphor* for a production process; the goal is to see to what extent the structure and operation of the metaphoric system matches up with what we already know of an industrial-production process, and even more importantly, at what point the production process radically departs from the structure and behavior of the known system. The point here, of course, is to choose the known system in such a way that the important production-process issues discussed in Section III can be easily formulated. At this point, each of the metaphors outlined is really only an idea for an idea, but I believe that they all contain the potential for development into a modeling framework suitable for basing a theory of manufacturing upon. Later sections of the paper give some indications as to how this might be done. But for now let us turn to the metaphors themselves.

A. ELECTRICAL CIRCUITS

Upon first consideration of the production process as a system, the situation faced today by manufacturing engineers seems to bear a striking resemblance to that encountered by electrical-circuit designers in the late 1940s and early 1950s. In both cases the problem is to transform a given set of inputs into specified outputs by means of a set of processing elements in such a way that the processing elements and their interconnections are in some sense optimal. In the case of circuits, the inputs and outputs are specified voltage patterns and the processing elements are various passive devices such as resistors, capacitors, and inductors, as well as active elements such as transistors and diodes; for industrial production the inputs are physical materials like steel, plastic, glass, and rubber; the outputs are finished assemblies such as cars, TVs, and cameras, whereas the processing elements are machines that act upon the material by cutting, bending, drilling, and so forth.

At first sight this analogy between electrical circuits and industrial-production processes looks promising. Furthermore, if a meaningful correlation can be made it will enable us to employ the vast array of mathematical and modeling results developed over the last 30 years in the field that we now know as linear (and nonlinear) system theory. Some significant obstacles, however, lie in the path of using an electrical circuit as a model for a production process.

The most obvious difference between the circuit and the production system is the way in which the material flows through the processing stages. In the circuit, the material that flows through one processing element to another is totally homogeneous in form: electrons. Furthermore, for most circuits of everyday interest one can assume that the flow is virtually instantaneous. Consequently, to analyze the circuit and develop the mathematical machinery for its design and production, we need not take into account any time-lags for electrons moving from one processing stage to the next; such movement occurs more or less simultaneously. On the other hand, in an industrial-production setting, neither of the foregoing simplifications can be accepted.

To begin with, in an industrial process the material flowing through the system is highly heterogeneous and, in contrast to circuits, not every processing element can accept all of the material types flowing through the system. So a theory of production processes will have to distinguish among the many incommensurate types of materials flowing through the system. At the bare minimum, this will require thinking of each of the communication channels in the process as consisting of many distinct physical conduits, one for each type of material, in contrast to the electrical circuit in which a wire from one element to another is just a single channel.

In addition, for production processes, time-lags count. Consequently, any theory of such processes will have to *explicitly* account for the delays involved in moving materials from one processing element to another. Within the context of a distributed manufacturing system, these time-lags can be expected to take on even greater significance, because the physical distances from one processing station to another may be orders of magnitude greater than they were in earlier, single-location plants.

Unfortunately, the problems of heterogeneous material flow and time-lags are not the only ones that stand in the way of using an electrical circuit as a mathematical paradigm for a theory of modern production processes. A significant feature of the factory of the future will be its level of automation. This implies that the factory has a capability for self-diagnosis, automatic reconfiguration, self-repair and, perhaps, even replication. Even modern mathematical models of electrical circuits do not contain explicit provision for these properties, so we must seek an extended paradigm in the area of living systems in order to accommodate these features within a model for production systems.

B. METABOLISM-REPAIR (M,R) SYSTEMS

The circuit-theory paradigm for production processes involves a diagram of the form

$$\Omega \xrightarrow{f} \Gamma ,$$

where Ω is a set of admissible inputs (labor, materials, money, equipment) and Γ is the set of output products. The map f is the rule, or prescription, by which the input quantities are operated upon by the production system to produce the desired products. So, in essence the entire production scheme is encapsulated in the behavioral map f. For circuits, f is physically realized by a collection of elements (resistors, capacitors), together with a wiring diagram specifying how the elements are to be connected to produce the correct circuit responses; in a manufacturing setting, the embodiment of f consists of various types of processing elements, usually machines of different sorts, together with the assembly plan by which the machines operate upon the input materials from Ω. As already noted, this picture is incomplete in a variety of ways as a model for actual manufacturing processes. It is especially deficient in that it contains no ready means to account for repair and replication, which are essential features of any truly automated production system. In order to address these shortcomings, we can think of the production process as a *living* system and can consider a metaphor based upon cellular metabolism.

Because any good theory of manufacturing must be able to account for operation concerns of day-to-day production and planning—items like costs, capacity utilization, inventories and the like—let us first examine the ways in which various aspects of cellular organization can act as counterparts of such features of manufacturing systems.

1. *Direct labor cost*. Almost all cellular processes are energy driven. A natural measure of the cost of a process is the number of high-energy phosphate bonds, or molecules of ATP, required to carry out the process. If desired, this can be expressed in more physical units of kcal/mole. Such units, which can play the same role in the accounting of cellular processes that money plays in economics, seem a natural analog of direct costs.
2. *Indirect labor cost*. An important part of the cellular plant is devoted to actually producing ATP from ingested carbon sources (e.g., sugars or other foods). ATP is needed to drive the cycles that make ATP, and this might be the natural analog of overhead or other types of indirect costs.
3. *Material costs*. Most of the raw materials required by a cell cannot, by themselves, get through the cellular membranes. They must be carried across by specific machinery ("permeases"), which must be manufactured by the cell. The cost to the cell of this manufacture and deployment of permease systems seems the most natural analog of the material costs of a human manufacturing system.
4. *Inventory cost*. Cells do not generally maintain pools of unutilized materials or inventories. Rather, they seem organized around a "just-in-time" strategy.
5. *Quality cost: Prevention, internal, external*. Cells have evolved a broad spectrum of quality controls. There are, for instance, a variety of error-correcting mechanisms and error-prevention mechanisms. An example of the latter are catalase and peroxidase molecules to deactivate potent oxidizing agents arising as the result of normal metabolism or from the effects of free radicals. A family of DNA repair enzymes exists to resplice strands of DNA broken by ionizing radiation or other mechanisms. The whole DNA-protein coding machinery is, in a precise sense, error correcting to protect itself against translation errors. If a defective protein or metabolite is made, digestive enzymes, either free or bound in organelles called lysosomes, take them apart. Of course, these diverse quality-control mechanisms cost something (in ATP), which seems the homolog of manufacturing quality costs.
6. *Reliability of output capacity*. Insofar as this relates to down time of manufacturing machinery or absentee personnel, this can be subsumed under the heading of quality control.
7. *Capacity utilization*. Just as in the human situation, the utilization of the products of cellular metabolism takes place outside the cellular system itself. For instance, if the cell is part of a multicellular organism and manufactures a hormone for utilization elsewhere, then the "market" for the hormone depends on exogenous circumstances. In fact, in such situations control mechanisms exist to adjust the capacity to the "market."
8. *Make-or-buy decisions*. This is an interesting aspect of cell biology. It bears on the notion of "flexibility," which in biology is often called plasticity or adaptability. It involves the situation in which an essential metabolite, like an amino acid, can be manufactured by the cell but also imported from the environment if it is available. If the metabolite is present in the environment, the entire metabolic pathway leading to its synthesis is shut down or inhibited and will remain so for as long as exogenous metabolite is present. Although no calculations have been performed, we can suppose that for cells we could show that it was cheaper (in ATP) to bring in exogenous metabolite than to synthesize it directly in cases in which both capabilities are present. In addition to these fast-acting inhibitory controls at the pathway level, slower, long-term controls exist that actually repress the synthesis of the enzymes in the pathway. This combination of fast and slow controls provides the plasticity to deal with both transient environmental circumstances and longer-term trends.
9. *Flexibility to shift in product mix*. The cellular homologs here seem to be closely related to the control mechanisms we have discussed above, in which environmental circumstances of one kind or another modify both the cellular machinery itself and the relative rates at which the machinery operates.
10. *Flexibility to technology change*. The closest cellular homolog to technology change seems to be either spontaneous mutation or the invasion of a cell by some foreign DNA (viral or otherwise). Alternatively, in the cells of multicellular organisms, homologs may be found in the process of cellular differentiation, in which banks of genes that were previously inactive (repressed) are turned on. In this way, the cell receives the effect of a change in genome, and a corresponding overhaul of the entire metabolic machinery, without losing the property of being a cell.

To employ an evocative terminology, let us call the input-output behavior of our production system its "metabolism," and term f the *metabolic map* of the system. Thus, when everything is operating correctly, the production process "metabolizes" inputs $\omega \in \Omega$ according to the scheme f and produces a product $\gamma \in \Gamma$. Note that, in general, the system can produce many different products and if we change ω to some different input ω^*, then the *same* scheme f will produce a product $\gamma^* \neq \gamma$. On the other hand, using ω^* it may be possible to find a scheme f^* such that $f^*(\omega^*) = \gamma$; that is, we can offset a change in the inputs by a corresponding change in the system's metabolism in order to leave the final product the same. This idea has some bearing upon the notion of "flexibility" in manufacturing and we shall return to it below. For now, let us assume that the objective of the production process is to maintain the scheme f and that due to various external or internal disturbances (machine failures, external economic fluctuations, communication noise, and the like), the production process departs from the scheme f and we want to repair the system and restore its design metabolism. How can we naturally extend the modeling framework given earlier for electrical circuits to include a repair mechanism?

The key to answering this question is to recognize that somehow the system must utilize a part of its metabolic output γ to generate the mechanism needed to restore the scheme f. Of course, a small technical point here is that if the purpose of the production process is to produce, say, TV sets, then γ is a TV set of a certain type and imagining that we can take a few of these TV sets and somehow use *them* to create a machine that will repair the TV production scheme f is nonsense. We can easily skirt this difficulty, however, by imagining that we first transform a few TV sets into their monetary equivalents, and then use this capital to build and maintain the needed repair operations. So, in what follows, we shall not distinguish between the actual physical outputs in Γ and their monetary equivalents, using each interchangeably as circumstances dictate. Returning to the repair operation, because its purpose is to regenerate the metabolism f from the system output, we can denote this operation by the map

$$\phi_f : \Gamma \rightarrow H(\Omega , \Gamma) ,$$

where $H(\Omega , \Gamma)$ is the set of all possible metabolisms; that is, all possible production schemes. Here we explicitly display the dependence of the repair map upon the metabolism f, because the objective of the repair process is to restore f in the face of various disturbances suffered by the production system. Thus, we end up with the following diagram of a manufacturing process as a "metabolism-repair (M , R) system,"

$$\Omega \xrightarrow{f} \Gamma \xrightarrow{\phi_f} H(\Omega , \Gamma).$$

Roughly speaking, in biological terms, f corresponds to the production system's "phenotype," whereas ϕ_f corresponds to the processes's "genetic" capacity to repair metabolic disturbances arising during the course of the system's operation and interaction with the outside world.

To complete the metaphor of the production system as a biological organism, we must address the issue of how to repair the repairers. The repair mechanisms were introduced to account for the fact that during the course of time, the metabolic machinery will erode and decay, thereby requiring some sort of rejuvenation if the production system is to avoid extinction. Precisely the same argument applies to the repair mechanism, but introducing repairers for the repairers and so forth, ending up in a useless infinite regress, is not helpful. The way out of this loop is to make the repair components self-replicating. In this way, new copies of the repair mechanism are continually being produced, and it is unnecessary to assume the repair functions are immortal or fall into an infinite regress of repairers to insure survivability of the process. Let us see how to introduce the idea of replication into the foregoing framework.

Because the replication operation involves reproducing the genetic component ϕ_f from the metabolic activity, the replication map, call it β_γ, must be such that

$$\beta_\gamma : H(\Omega , \Gamma) \rightarrow H[\Gamma , H(\Omega , \Gamma)],$$

if it exists at all. The question is: How can such a map β_γ be constructed from the basic metabolic components Ω , Γ , and $H(\Omega , \Gamma)$? To see how this is done, consider a somewhat more general situation.

Let X and Y be arbitrary sets. Then for each $x \in X$, we can define a map

$$\hat{x} : H(X , Y) \rightarrow Y$$

by the rule

$$\hat{x}(f) = f(x),$$

for all $f \in H(X , Y)$. Thus we have an *embedding* of X into the set $H[H(X , Y) , Y]$. Now, *assume* that the map $\hat{x}$ has a left inverse $\hat{x}^{-1}$, so that,

$$\hat{x}^{-1} : Y \rightarrow H(X , Y).$$

Then we clearly have

$$\hat{x}^{-1} \hat{x}(f) = f$$

for all $f \in H(X , Y)$.

Returning now to our replication situation, we set

$$X \doteq \Gamma, \quad Y = H(\Omega , \Gamma),$$

and apply the foregoing general argument to obtain for each $\gamma \in \Gamma$, a map $\beta_\gamma \doteq \hat{\gamma}^{-1}$ with the property that

$$\beta_\gamma : H(\Omega , \Gamma) \rightarrow H[\Gamma , H(\Omega , \Gamma)]$$

for all $\hat{\gamma}$ possessing a left inverse. In short, the metabolic activity can be used to reproduce its repair component if the technical condition on the invertibility of the map $\hat{\gamma}$ is satisfied. The manufacturing interpretation of this condition is that $\hat{\gamma}$ is invertible if different repair processes (i.e., different genetic mechanisms $\phi_{f_1} \neq \phi_{f_2}$) give rise to different metabolic processes $f_1 \neq f_2$.

The minimal structure introduced thus far to define an (M,R) system is already sufficient to shed light on a variety of interesting questions surrounding the ways in which a firm can respond to changes in its operating environment, the possibility for innovation to occur through environmental effect, the circumstances under which environmental changes can be reversed, feedback, and so on. Here we sketch the way in which these issues appear within the (M,R) framework and consider the conclusions that can be drawn about production-system behavior from this structure.

1. Stable Metabolic Operations in Changing Environments

Imagine the situation in which the "usual" input ω of raw materials, labor, and so forth is disturbed to a new input $\overline{\omega}$. The condition for stable operation of the production process is for the environment ω to be such that

$$\phi_f [f(\omega)] = f , \qquad (*)$$

that is, the metabolic structure f is stable in the environment ω in the sense that the repair mechanism ϕ_f always regenerates f when the environmental input is ω. We would say that all $\omega \in \Omega$ satisfying (*) form a stable environment for the process.

Now suppose that the new environment $\overline{\omega} \neq \omega$. Then (*) will hold only if either

$$f(\omega) = f(\overline{\omega})$$

or

$$\phi_f[f(\bar{\omega})] = f.$$

The first case is trivial in the sense that the observed revenues of the firm are invariant to the change of environmental inputs. If $f(\omega) \neq f(\bar{\omega})$, then the outputs are not invariant with respect to the change of inputs and we must consider the repair mechanism to see whether or not the input alterations can be compensated for in the sense that

$$\phi_f[f(\bar{\omega})] = \bar{f} \neq f,$$

with $\bar{f}(\bar{\omega}) = f(\omega)$; that is, will the genetic mechanism produce a new metabolism $\bar{f}$ that will duplicate the outputs of f, but now using the input $\bar{\omega}$ rather than ω? In this case, the entire metabolic activity of the firm would be permanently altered if we had

$$\phi_f[\bar{f}(\bar{\omega})] = \bar{f}.$$

On the other hand, if we had $\bar{f}(\bar{\omega}) = f(\omega)$ or, more generally,

$$\phi_f[\bar{f}(\bar{\omega})] = f,$$

then the metabolism would only undergo periodic changes in time.

Finally, we could have the situation in which

$$\phi_f[\bar{f}(\bar{\omega})] = \hat{f} \neq f, \bar{f},$$

and, iterating this process, we see that an input change may cause the production process to wander about in the set $H(\Omega, \Gamma)$, changing its operating scheme through a sequence of metabolic processes $f^{(1)}, f^{(2)}, f^{(3)}, \ldots$. This "hunting" process will terminate if either

1) an N exists such that

$$\phi_f[f^{(N)}(\bar{\omega})] = f^{(N)}$$

or

2) an N exists such that

$$\phi_f[f^{(N)}(\bar{\omega})] = f^{(N-k)}, \qquad k = 1, 2, \ldots, N-1.$$

In case 1) the process becomes stable in the new environment $\bar{\omega}$, whereas in case 2) the process undergoes periodic changes in its metabolic structure. If no such N exists, then the process is unstable and aperiodic. (Note: This last possibility can occur only if the set $H(\Omega, \Gamma)$ of possible production procedures is infinite.)

2. Lamarckian Changes in the Repair Process

The above discussion of metabolic changes was undertaken subject to the tacit assumption that the repair map ϕ_f remains unchanged. We are interested in inquiring whether or not an environmental change $\omega \to \bar{\omega}$ can generate a Lamarckian type of genetic change in ϕ_f through the replication process described earlier. If such a change is possible, then we could infer that the actual repair process that regenerates the metabolic activity f could be affected by environmental changes alone.

To examine this question, suppose we have the environmental change $\omega \to \bar{\omega}$. Then the replication map β_γ associated with the input ω and the output $\gamma = f(\omega)$ is changed to $\beta_{\bar{\gamma}}$, where $\bar{\gamma} = f(\bar{\omega})$. Recalling that

$$\beta_\gamma = \hat{\gamma}^{-1}, \quad \beta_{\hat{\bar{\gamma}}} = \hat{\bar{\gamma}}^{-1}$$

and

$$\hat{\gamma}(\phi_f) = \phi_f[f(\omega)], \quad \hat{\bar{\gamma}}(\phi_f) = \phi_f[f(\bar{\omega})],$$

after applying β_γ, $\beta_{\hat{\bar{\gamma}}}$, respectively to the last two relations, we find that

$$\beta_\gamma\{\phi_f[f(\omega)]\} = \beta_{\bar{\gamma}}\{\phi_f[f(\bar{\omega})]\} = \phi_f,$$

showing that the new replication map $\beta_{\bar{\gamma}}$ replicates the existing repair component ϕ_f exactly. Thus, an environmental change alone can have no effect upon the repair map ϕ_f.

Now we ask whether it is possible for a change in the metabolic production procedure to result in a change of the firm's "genotype". Suppose we replace the metabolic activity f with some other production process $b \in H(\Omega, \Gamma)$. By definition

$$\hat{b}(\omega)(\phi_f) = \phi_b[b(\omega)].$$

Assuming $\hat{b}(\omega)$ is invertible, we apply $\hat{b}^{-1}(\omega)$ to both sides of the above relation to obtain

$$\beta_{b(\omega)}\{\phi_f[b(\omega)]\} = \phi_f.$$

Thus, the induced replication map reproduces the original repair component of the firm under all conditions. In short, no Lamarckian changes in the metabolic component, either in the environment Ω or in the metabolic set $H(\Omega, \Gamma)$, can result in changes in the repair mechanism. Such "genetic" changes can only come about through a direct intervention in the genetic code itself (mutation) and not via indirect metabolic alterations.

3. Feedback as an Environmental Regulator

The environmental changes discussed thus far have been assumed to be generated by actions external to the production process itself. The output of revenue, however, may often be employed as one of the environmental inputs; that is, ω is a function of γ, $\omega = \omega(\gamma)$. In this event, the process actually creates part of its own environment and, as a consequence, can partially regulate its own structural alterations. An important aspect of this general process is to understand to what degree adverse environmental disturbances can be "neutralized" by suitably chosen feedback policies. This question is a special case of the more general problem of "reachability," in which we ask about the possibility of attaining any predefined metabolic structure by means of a sequence of environmental changes.

Up to this point, we have assumed that the metabolic and repair functions take place instantaneously; that is, inputs are transformed into products immediately, and no delay occurs in either repairing the metabolic process itself or in the replication of the repair mechanism. Needless to say, these assumptions are pure fiction; production and repair/replication take time.

Although we have no space here to enter into a detailed discussion of the matter, let us simplify the situation by assuming only two types of delays. The first we term the *production delay*, corresponding to the time required to transform a given input of materials, manpower, and knowledge into an observable product or the time required for a repair function to restore a metabolic operation. The second type of delay we shall call the *transport delay*. It corresponds to the time needed to transport the output to a place at which it can be utilized as the input to either another process or to a repair mechanism (or to the external world).

Closely related to the time-lag problem is the matter of system dynamics. Several issues surround this topic, not all of them mutually consistent. For simplicity, let us consider here only the case of a single process modeled as an (M,R) system. Abstractly, the diagram for the process is

$$\begin{array}{l} \Omega \xrightarrow{f} \Gamma \rightarrow \Phi_f \rightarrow H(\Omega\,,\Gamma) \xrightarrow{\beta_\gamma} H[\Gamma\,, H(\Omega\,,\Gamma)] \\ {}_{g}\searrow \; \nearrow^{h} \\ \quad X \end{array}$$

and ask in what manner the process can be regarded as a dynamic system. If it were not for the repair and replication maps Φ_f and β_γ, this would be a straightforward question addressable by normal system-theoretic realization theory procedures; that is, we would have the problem of constructing a canonical internal model of the process

$$\dot{x} = p(x\,,u),$$
$$y = h(x),$$

whose input-output behavior duplicates that of the given metabolic map f. Techniques for handling this question are readily available in the mathematical systems-theory literature [1, 2, 3].

Let us ignore for the moment the factorization of the metabolism f through maps g and h, and consider the (M,R) system

$$\Omega \xrightarrow{f} \Gamma \xrightarrow{\Phi_f} H(\Omega\,,\Gamma)$$

where $f \in H(\Omega\,,\Gamma)$, $\Phi_f \in H(\Gamma\,, H(\Omega\,,\Gamma))$. We wish to show how this abstract model of a production process can be considered as a sequential machine; that is, as a discrete-time dynamic input-output system.

Let us recall that a sequential machine M is a composite $M = (A\,,B\,,S\,,\delta\,,\lambda)$, where A, B, and S are sets (possibly infinite), whereas $\delta : A \times S \rightarrow S\,, \lambda : S \times A \rightarrow B$ are maps. We interpret A as the input alphabet of M, B as the output set, S as the set of states, with δ and λ the state-transition and output maps of the machine, respectively. At each discrete instant of time $t = 0\,,1\,,2\,,\ldots\,,M$ receives an input symbol from A, emits an output in B, and the state is changed according to the rule δ, and the process continues from the time $t + 1$. Further details on the properties of sequential machines can be found, for example in [2, 3].

In order to characterize the production process as a sequential machine, we make the identifications

$$A = \Omega, \quad B = \Gamma, \quad S = H(\Omega\,,\Gamma),$$
$$\delta(\omega\,,f) = \Phi_f[f(\omega)], \quad \lambda(f\,,\omega) = f(\omega).$$

Thus, in general, any process can be regarded as a sequential machine in which the set of "states" of the machine correspond to the set of possible "phenotypes" of the process, whereas the input and output sets of the machine are the inputs and outputs of the process, respectively.

Putting the above ideas together, we arrive at the following scheme for characterizing the dynamics of the process:

1. Regarding the process as a sequential machine formed from the elements $F = [\Omega\,,\Gamma\,,H(\Omega\,,\Gamma), \Phi_f, f]$, we compute the metabolic process f at the time $t = 0$;
2. Using f, Ω, Γ, we employ realization theory to form a canonical model for the state space X and the maps g and h;
3. Let $t \rightarrow t + 1$ and use the sequential machine to calculate the new metabolism f. If it is the same as at the previous time step, then continue to use the earlier X, g, and h; if f changes, calculate a new canonical model and continue the process with the new model until the next time period.

In the above scheme, a change of metabolism implies that the production process has been changed. This can come about only if the repair map Φ_f fails to reproduce f. We have already seen that this may come about only by means of environmental changes, in general, unless the replication map fails to exist. But this last situation depends entirely upon the size of the set $H(\Omega\,,H(\Omega\,,\Gamma))$, the space of all possible repair maps. If it is either too large or too small, then no replication is possible. It would take us too far afield to enter into the details of this argument here, but the implications are that only in a highly restricted class of categories can replicating (M,R) systems exist, and we must search within this class for viable models of industrial production.

Before leaving the (M,R) system metaphor, note that the treatment given above is purely *functional;* that is, the framework specifies the functional role of the metabolic map f, the repair process Φ_f, and the replication operation β_γ, but it says nothing about the operational details of the ways in which these functions are to be carried out. Such details are of the essence in transforming the above theoretic framework into an effective theory of manufacturing processes. As an example of the type of research that will be required in this direction, the repair map Φ_f represents the "genotype" of the production system. To specify the detailed operations of this function, we must develop an industrial analog of the "genetic code" and specify in detail the "DNA" of the firm. At first glance, this may appear to be stretching the biological metaphor beyond its elastic limit, but when we recall the discussion given earlier concerning formal systems and industrial processes and then examine the direct linkages between cellular genetics and formal systems as given, for example, in [4], then the whole idea takes on a much more plausible tone. This entire circle of ideas strongly suggests our next metaphor for production, which involves consideration of the production system as a *language*.

C. NATURAL AND ARTIFICIAL LANGUAGES

A very speculative, but nevertheless intriguing, metaphor for production processes is offered by consideration of such a process as a language. To see what's involved in setting up such a framework, consider a simplified language consisting only of nouns, verbs, adjectives, adverbs, and conjunctions. We can then make associations between these elements of the language and the functional elements of a production process as shown in Table 1.

The rules of syntax tell us how language elements can be combined to form *grammatically correct* sentences, and the theory of semantics enables us to ascertain the *meaning* of admissible sentences. Using the associations of Table 1, we would interpret the "syntax of production" to be the rules that tell us in what sequence various materials could be processed through the production system, and the "semantics of production" would enable us to distinguish products that perform a functionally meaningful role from those that do not. For example, the grammatically correct, but nonsense, statement *John and air* might correspond to the physically realizable, but useless, product consisting of a steel plate (John) joined to a fountain pen (air), say. On the other hand, using the same identification, the meaningful statement *John and Jim* may correspond to the useful product of a door, consisting of the steel plate joined to a doorknob (Jim). By forming more elaborate sentences, including more language elements and making finer associations between the language and production elements, one is tempted to speculate that every meaningful sentence corresponds to a production process that outputs a functionally useful product. Because we know that natural languages are highly redundant and admit the possibility of expressing the same thought in many different ways, if the foregoing conjecture could be supported, then we would then possess the basis for generating a myriad of alternative production processes all producing the same final product. Let us consider an example of how such an isomorphism would work.

As before, let John represent a steel plate and suppose that we have one machine that *drills* a one half–inch hole in such a plate, and a second machine that *punches* a similar hole. Call the first machine operation "eats", and call the operation of the second machine "consumes." Further, the physical object consisting of a steel plate with a one half–inch hole in it is labeled "bread." Then, with these identifications, the subject-verb-object sentences *John eats bread* and *John consumes bread* are totally equivalent semantically, and furthermore they represent alternate production schemes that produce an identical final product. The first scheme drills the hole in the plate, and the second punches it, with the transformation of John resulting in the same end product. Note also that the syntactically correct, but semantically meaningless, sentences *The bread eats John* and *The bread consumes John* are meaningless from a production point of view as well, because both imply that a steel plate with a hole in it when processed by either the drill or the punch will result in a plate with no hole.

The preceding ideas can be greatly extended and enhanced by employment of Chomsky's theory of generative transformational grammars [5]. A generative grammar is a finite system that contains explicit and formal descriptive rules that assign to the sentences of the language correct descriptions of their structure and associates their form with a representation of their meaning. Thus, a *generative* grammar is an algebraic system of

TABLE 1
Language and Production Analogs

Language Element	Production Element
Noun	Physical unit of material
Verb	Processing (transforming) operation
Adjective	Material characteristics
Adverb	Modulation of processing operation
Conjunction	Joining/separating operation

formal rules that can assign the correct structural descriptions to an infinite number of well-formed sentences in a language. To say that a grammar is *transformational* means that it contains a certain formal device, namely a rule for moving an element of a sentence elsewhere in particular prescribed ways. The overall structure of such a grammar is depicted in Figure 3.

To see how this system works, we can use a scheme similar to a decision tree, called a *phrase marker,* in order to first divide a sentence into its syntactical parts. To make things simple, consider the sentence *John must eat spinach.* Starting at the left, we have the phrase structure rules (PS) and the lexicon. There is only one set of rules for each grammar and they define the structural relations among the various units. The lexicon specifies properties of lexical items; for example, that *eat* is a verb (V) that may be followed by a direct object noun phrase (NP). The initial phrase marker generated by the PS rules and lexicon for the above sentence is depicted in Figure 4. (In Figure 4, *Spec* denotes a specifier of the verb phrase, and the bar denotes a dominance relation.)

The transformational rules provide procedures for moving constituents to different places in the phrase marker. For example, the rule $(\cdot\cdot\cdot)\text{AUX} \rightarrow \text{AUX}(\cdot\cdot\cdot)$, where $(\cdot\cdot\cdot)$ stands for anything else in the phrase marker, has the effect of moving the auxiliary item to the front of everything else. Applying this rule to the initial phrase marker above yields the new sentence *Must John each spinach.* These two components, the base and the transformation, make up the syntactic component of a language. The output of this component is the so-called *surface structure,* which serves as the input to both the phonetic and semantic interpretation. It would take us too far afield to enter into a consideration of the semantic component here, but it should be noted that any attempt to employ the foregoing language metaphor for manufacturing will have to confront both the syntactic and semantic structure of the production process.

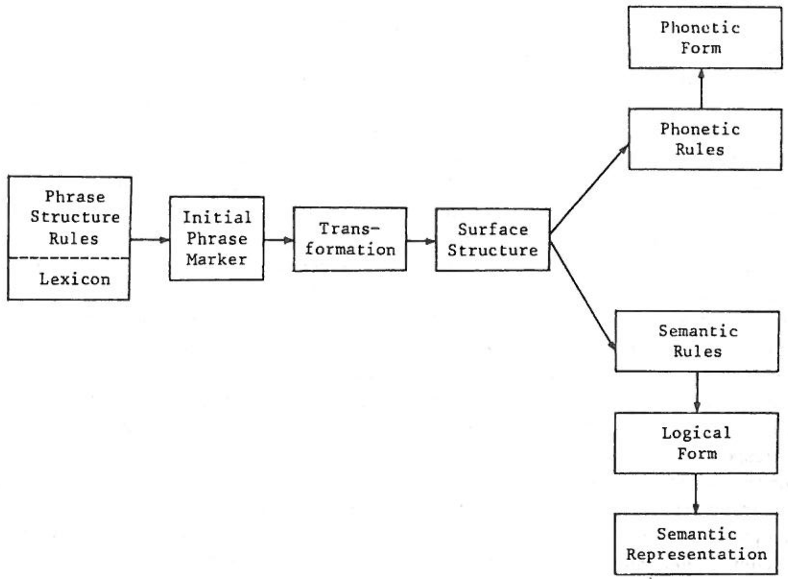

Fig. 3. General form of a generative transformational grammar.

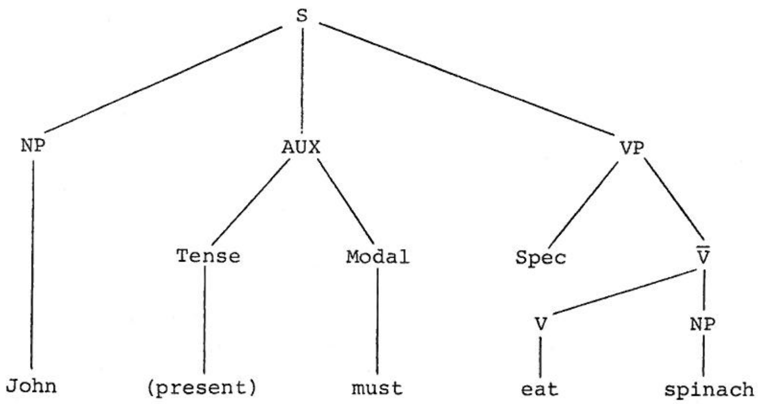

Fig. 4. Initial phrase marker.

In view of the elementary example given earlier, the relevance of Chomsky's linguistic program for a general theory of manufacturing processes is evident. Not so evident and requiring a systematic research program to be clarified is whether or not the full machinery of generative transformational grammars as developed for human languages is really needed for the vastly more restricted domain of industrial manufacturing. After all, the range of "expressions" that a language of manufacturing needs to convey is far more restricted than is that for any natural human language, so perhaps a theory of manufacturing could be based upon either a more elementary type of generative grammar or a more restrictive artificial language.

Originally, Chomsky tried to develop his idea of a generative grammar using *finite-state* models. In such finite-state grammars, sentences are generated by a series of choices made one after the other. First one word is chosen, then the second, and so on. Using Markov processes, each choice is determined by probabilities that are a function of preceding choices. The rules defining the transition from one word to the next are given in terms of probabilities. Although such a scheme is unsuitable for human languages because such systems cannot deal with nonadjacent dependencies, the finite-state grammars, suitably modified, may at least provide the basis for a linguistic theory of production.

In a seemingly different direction, we can also consider the use of an artificial language as a model for manufacturing. Probably the most straightforward approach is to consider a *computer language,* such as LISP or BASIC. Computer languages have the great advantage, in comparison with natural languages, of a very limited lexicon as well as a totally unambiguous and very streamlined grammar. In fact, for such languages the finite-state grammars can serve as a basis for construction of a theory of production. For example, in the LISP language the basic elements are lists of atoms, together with functions that create, destroy, and modify these lists. In the manufacturing context, the atoms would correspond to various individual parts of an object and the *lists* would be assemblies of subcomponents, whereas the *functions* would correspond to operations performed on the subassemblies to transform them into the "ultimate" list corresponding to the complete end product. The grammar then tells us which operations can be performed upon which assemblies and what new assemblies will result. It is then a problem of "reachability" to decide whether or not the desired end product can be constructed using the raw materials and operations available and, if so, in how many ways.

Before leaving the topic of language as a metaphor for production, note that most of what has been said above pertains to the syntactic component of Figure 3; that is, the surface structure, with no attention given to the semantic aspect of language. According to most contemporary linguistic theorists, the determination of the *meaning* of a sentence requires another set of rules, termed *interpretive rules,* that for any given surface structure derives a representation of the sentence called its *logical form.* The logical form then specifies the scope of quantifiers, relations among noun phrases, and so on, which give meaning to the sentence. A detailed analysis of semantics, as well as syntactics, would be required to use the linguistic framework as a basis for a theory of production because we must be able to specify exactly what function is served by the item being produced. Finally, we note a very provocative theory of semantics that has been put forward by the topologist R. Thom [6]. In this theory, Thom explores the relationship of spatiotemporal events, their representation as sentences in a language, and their topological character as structurally stable forms. In this manner, he is able to attach a well-defined archetypical morphology to each such event and show how one morphology gives way to another during the course of a sentence describing an event. It is tempting to speculate that incorporation of some of these concepts may prove expedient for a linguistic theory of production, which deals only with sequences of spatiotemporal events.

D. CELLULAR AUTOMATA, COMPUTER NETWORKS, AND PARTICLE PROCESSES

In addition to the candidate theoretical structures for a theory of production given by life, languages, and electrical circuits, quite a number of alternative metaphors also seem to be plausible contenders. Because these metaphors have not yet been extensively explored in the manufacturing context, we shall only briefly describe here the basic idea underlying them, deferring a more detailed treatment for another occasion.

To begin with, the idea of using a computer language (software) as a metaphor suggests the immediate question: why not use hardware instead? Why not, indeed! In fact, a natural approach is to think of the production process as a special type of *cellular automaton.* Each cell in the grid of the automaton corresponds to a processing element of the production system, with the rule of operation of the cell corresponding to the type of operation performed by the production element upon its input materials. The "wiring diagram" of the automaton would then be specified by the way in which the production system can be sequenced, with the output of each cell corresponding to the physical output of the production element represented by the cell. Because a voluminous (and ever-expanding) literature already exists outlining the properties of such automata and their use as models for everything from DNA sequences to galactic clusters [7], we certainly appear to have ample opportunity to exploit this generality in the context of a theory of manufacturing. The biological metaphor presented earlier has a very strong connection with automata-theoretic models as is evidenced, for example, in [8]. In a similar vein, a close connection exists between automata and languages, as outlined in [9, 10]. These formal relationships provide the basis for a unified consideration of a number of important manufacturing concepts such as complexity and flexibility, as will be discussed in the following section.

We may also consider a computer system as a metaphor for manufacturing in another manner. A production system is fundamentally a movement of *materials* from one processing station to another; a computer system moves *information* from one station to another. Making the obvious connection between information and materials, we can think of a modern parallel-processing computer network as an information-processing analog of a production process. This metaphor is rather appealing because many of the same features that characterize industrial production (incommensurate materials, time delays, distributed processing), also enter into the operation of a computer network. Furthermore, the classic von Neumann–type serial-processing architecture is now undergoing major transformations with the current "new wave" architectures modeled closely upon the cellular automata structures discussed above [11].

Finally, let us note a very speculative metaphor for manufacturing built around the idea of a particle process. In elementary particle physics, the essential events are the collision of one particle with another to form a shower of new particles. In producing a given industrial product, somewhat the same kind of events occur: one physical item, a subassembly say, "collides" with another item with the result being that the original items disappear and a new "particle" emerges.

Recently, a very powerful approach to the modeling of particle processes in physics by abstract computational processes has been developed in [12]. The essence of this theory is the coding of a particle by a *message* consisting of a finite sequence of bits. A message is processed by one of a set of computational actors that can change at most one bit in the message. In the theory, only six types of actors occur: source, sink, decider, arbiter, synchronizer, and memory. The details of each actor's operations are not important for now; we need only note that the concept of a collision of two particles corresponds to a bit exchange through a common memory actor by two messages.

For manufacturing processes, it seems plausible that the same abstract computational metaphor can be applied by coding each distinguishable type of part and/or subassembly by a bit string and regarding the primitive actors above as abstract analogs to distinct production operations. Note that the set of actors given in [12] is not the only set that will work, the only requirements being:

1. Each actor examines or changes at most one bit of a message;
2. No actor's operation can be expressed in terms of the others;
3. All effective computations can be modeled; that is, the set of actors has the power of a Turing machine;
4. At each moment, an actor is either processing a single message or is idle; that is, each actor processes messages serially.

With these requirements, we can show that the computational metaphor enables one to reformulate much of quantum particle physics in simpler terms and, for example, to see how one can retain the idea of cause and effect even in the face of indeterminism. Our question is whether or not one can formulate manufacturing processes in computational terms. The hunch here is that we can, but it is a hunch that will require a serious research effort to substantiate.

E. CHEMICAL ENGINEERING AND REACTION METAPHORS

Note that many of the features cited as being characteristics of the factory of the future—continuous operation, smooth transition over a wide range of raw material inputs and market demand outputs, self-diagnosis and repair—have been present in the chemical-engineering world for a number of years. These features emerged as a consequence of the development of flow processes and the unit operations approach that came out of the universities and not the industry itself. Many differences exist between the continuous-flow processes prevalent in chemical engineering operations and the discrete-product, batch processes found in most industrial processes. Nonetheless, speculation that some of the ideas that make the system work for chemical processes can be tailored for use in other areas is tempting.

In a quite different direction, we can also speculate upon the notion of a chemical-reaction metaphor for manufacturing. A basic problem in most chemical experimentation is to arrange a sequence of viable and efficient reactions to create a given compound from a selection of more basic chemicals. We might think of the component chemicals as the assemblies of the final desired product, with the individual reactions and reaction rates representing the processing stages. Or perhaps external factors like heat, light, and mixing, together with catalysts of various sorts, would play the role of individual processing steps. Finally, the laws of chemical combinations would specify the types of operation that are theoretically feasible and the order in which they may be performed. Finally, such a metaphor is closely related to the linguistic scheme discussed earlier, although the chemical-reaction metaphor may be easier to deal with, because the elements and grammar are far more well understood than is the case for natural languages.

6. System Concepts and Manufacturing Metaphors

Earlier we briefly mentioned a number of system-theoretic concepts—flexibility, complexity, reliability, . . . that enter into the study of manufacturing processes. With the preceding modeling metaphors as background, let us now consider a few of these conceptual topics in somewhat greater detail.

A. RELIABILITY

An essential feature of an autonomous and automatic manufacturing process is that it be able to monitor itself and effect needed repairs to components that depart from specifications. Such "self-correcting" abilities ensure what we usually think of as the *reliability* of the process.

Within the framework of (M,R) systems, we have already explicitly incorporated a "repair-and-replication" mechanism that functions to preserve the design purpose of the production process in the face of natural senescence (wear), as well as external disturbances. We have already noted that in order to explicitly spell out the workings of the repair component at the level of an *individual* production cell, it will be necessary to construct the industrial analog of DNA and its self-reproducing mechanisms. This is a problem for future research. But for now, let us accept the unreliability of individual-system components and consider the question of whether or not a *network* of production cells can operate reliably even though the individual components may fail.

In order to fix ideas, consider the specific (M,R) network depicted in Figure 5. The square blocks labeled $F_1, F_2, \ldots, F_6$ represent the metabolic processes of the individual cells, and the ovals denoted $R_1, R_2, \ldots, R_6$ represent the respective cells' repair mechanisms. The requirements that we impose for any such network are modest:

1. Each cell must receive at least one input, either from the external world or from the output of another cell;
2. Each cell produces at least one output; and
3. Each repair mechanism receives the output of at least one cell in the network.

The first general issue to consider for an (M,R) network is the dependency structure. We are concerned with the question of how the removal of a given cell from the network affects the existence and operation of other cells in the production process. For instance, referring to Figure 5, we see that the failure of cell F_5 results in the failure of cell F_6, as well, because F_6 receives its only input from F_5. Furthermore, the failure of F_5 *may* influence the operation of F_3, because F_3 receives part of its input from F_5. Thus, in this case we would consider cells F_3 through F_6 to make up the *dependency set* of cell F_5. Any cell the failure of which affects either the existence or operation of *all* cells in the

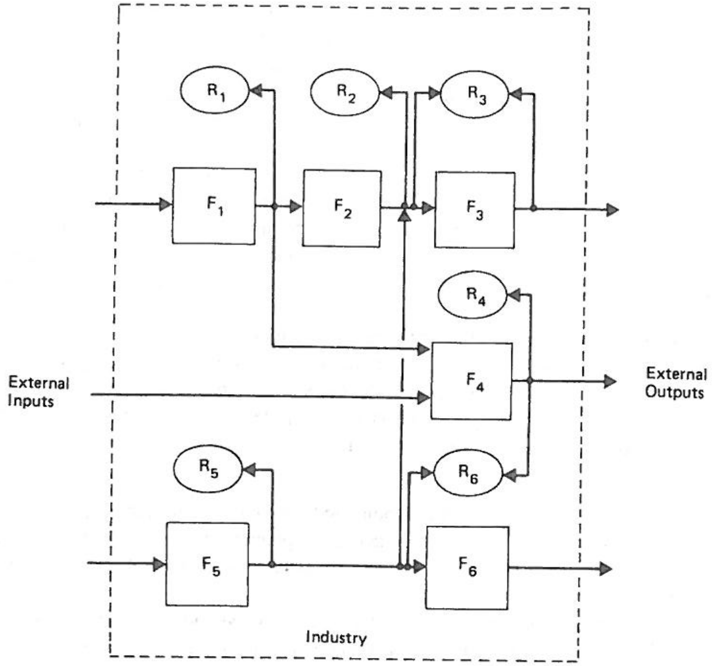

Fig. 5. A Typical (*M,R*) Network.

process will be called a *central* cell; that is, the dependency set of a central cell is the entire process.

Because we know that in the absence of the repair mechanism any cell will go out of existence after some finite lifetime, cells in the dependency set of a given cell will clearly also go out of existence when that cell does. With the repair mechanism in operation, however, a given cell quite possibly could "come back to life" even after its initial demise. For example, cell F_6 in Figure 5 may cease metabolic operation and be removed from the network; however, the repair mechanism R_6 receives its necessary inputs from cells F_4 and F_5, indicating that whatever "shock" caused the extinction of F_6, the cell will be re-inserted into the network after some characteristic delay time depending upon the repair mechanism R_6. In other words, copies of F_6 will continue to be manufactured even after the removal of F_6 from the network. Cells like F_6 will be called *reestablishable,* whereas all other cells are termed *nonreestablishable* (e.g., F_1 , F_5 , . . .). An important relationship exists between the notion of reestablishability and the concept of a central component expressed by the following result.

THEOREM 1. *Every* (M,R) *network must contain at least one nonreestablishable cell.*
COROLLARY. *If an* (M,R) *network contains only one nonreestablishable cell, then that cell is a central component.*

The significance of this result is twofold:

1. Every process must contain at least one cell the metabolic failure of which cannot be repaired. This conclusion follows only from the connective structure of the (*M,R*) network and is completely independent of the specific process, the procedures of the cells, their products, or production strategies. It is solely a consequence of the meaning of the metabolism-repair functions and the replication process.
2. In order to be "resilient" to unforeseen disturbances, one would desire a process to consist of a large number of reestablishable cells. On the other hand, the above results show that if only a small number of cells are nonreestablishable, then a high likelihood exists that one of them will be a central component the failure of which will destroy the entire process. Thus, a production process with a large number of reestablishable cells will be able to survive many types of shocks and surprises, but certain types of disturbances will effectively cripple the whole process. Consequently, a process with a relatively large number of non-reestablishable cells may be desirable if protecting the process from complete breakdown is desired.

The traditional engineering solution to the problem of system reliability is through redundancy. Several copies of critical system components are connected in parallel, with the failure of a component leading to a "back-up circuit" being placed in operation. Within the linguistic framework for production processes, reliability would correspond to a sentence structure that somehow preserves the semantic content of the sentence, even when key elements are either corrupted or removed entirely. For example, the almost subconsciously automatic correction of typographical errors, as when we interchange the *e* and *a* in *maesure* to recover the word *measure*. Such error-correcting mechanisms have been formalized in the limited context of information theory by the use of *error-correcting codes* and in computer-hardware communication by means of simple schemes like *bit-parity checks*. Such ideas represent different means to the same end: reliability through redundancy. The appropriate analogs for the production frameworks represent an important area for system-theoretic research in manufacturing.

B. FLEXIBILITY

The most significant distinguishing feature of the manufacturing system of the future is by almost unanimous agreement claimed to be its "flexibility." One of the aims of a *theory* of manufacturing is to provide a means with which the level of flexibility of different production processes can be compared. We now have available only the intuitive feeling that somehow flexibility means the ability of a fixed collection of production elements to be reconfigured, quickly and cheaply, to produce a large variety of different products. Superimposed upon this picture are additional intuitions involving small economic lot sizes, reduced inventory levels, high levels of automation and other features that are also supposed to reflect flexibility in some way. An initial step toward providing an *analytic* framework for discussing flexibility using graph-theoretic ideas is given in [13], but other than this recent work, the literature to date is remarkably silent on the overall systems view of this topic.

From the vantage point of the frameworks and metaphors we have outlined, a variety of directions can be explored to generate a consistent picture of what a flexible manufacturing system would be like. To begin with, we can consider a computation-process metaphor in which the flexibility of the system relates to the ability to carry out a variety of computations efficiently and quickly, using a limited number of computation actors suitably configured. This transcription of the intuitive notion of flexible manufacturing sounds suspiciously like a description of a general-purpose computing machine and this is not by accident! Little separates abstractly the general-purpose computer's ability to "simulate" any physically realizable process (the Turing-Church Thesis) and the concept of an *ideally flexible* manufacturing system: one that can be configured to produce any product that is physically producible (a universal constructor). Thus, if we employ a computational metaphor for manufacturing, a reasonable approach to characterizing a given production system's flexibility would be to ask, How far from a universal Turing machine is the system's computational model?

Within the context of (M,R) systems, the notion of flexibility can be stated in a quite different way. If flexibility is to mean anything, it must be related to the idea of a production system's ability to "metabolize" a specified output from a set of inputs without having to leave the structure of a self-repairing system. Stated another way, if we are given a particular desired production metabolism $g : X \rightarrow Y$, we ask under what circumstances g can be imbedded in an (M,R) system. This question can be approached through the notions of category theory, and its resolution depends critically upon the nature of the category to which g and the sets X and Y belong. The original production system, specified by the metabolism f with the input and output sets Ω and Γ, is then more or less *flexible*, as one can imbed a greater or fewer number of new metabolisms $g : X \rightarrow Y$ within the category formed by f, Ω, and Γ.

If we adopt a linguistic framework to characterize the production system, the problem of system flexibility revolves about the ability to generate a large number of *semantically distinct* expressions from a common *syntactic* structure. One way to attack this matter is by the use of *conceptual graphs*, enabling us to translate a given semantic content into a graph. Isomorphic graphs then contain the same semantic content and, hence, represent alternative ways of producing the same final product. Thus, the number of isomorphism classes associated with a given *set* of semantic networks gives a measure of the flexibility of the network to "express" concepts.

Let us now imagine the situation in which the structure of the semantic network is fixed. Further, assume that we have available a set of transformation $T = \{T_i\}$ that can be applied to the network such that each transformed network is syntactically admissible. The question is, To what degree do the transformed networks possess semantic content different from the original? The answer, of course, is intimately bound up with the structure of the transformation set T. For example, if T forms a group, then all transformed networks are equivalent and contain the same semantic message. Weaker structure on T admits the possibility of new meanings (products) emerging. The issue of production-system flexibility then comes down to a study of the properties of T.

C. COMPLEXITY

In the general system literature, complexity is a word that is almost as overworked as flexibility is in modern manufacturing parlance. Typically, complexity refers to some combination of a system's ability to manifest counterintuitive behavior, number of distinguishable components, density of subsystem interconnections, and so forth. As a general rule, complexity is something that is deemed to be bad, and means are sought for its characterization so that it can be reduced, presumably with the goal of returning a system to that state of bliss in which its manager can have a complete understanding of all its workings and be in a position to take whatever actions suit his or her perceived needs with no fear of hidden "time bombs" exploding unexpectedly. Whatever one may think of the dubious merits of such a viewpoint, the study of system complexity and its interrelationship with the other system concepts we are discussing contains a number of points of practical concern for production systems. For instance, if complexity in a production system means more components, then it is clearly of economic interest to isolate the "simplest" system that can produce a given item. Or, if a system is to be able to produce a fixed spectrum of products, how complex must it be? Problems of this type, and many more, force upon us a thorough investigation of the complexity concept.

From the outset one can take two qualitatively distinct views of the system-complexity issue: complexity is an *intrinsic* aspect of the system or it isn't! The latter view, based upon the idea that complexity is a *contingent* property of a system that comes about from the system's interaction with an observer, has been argued forcefully elsewhere [14], so here let us concentrate upon the "intrinsic school" of complexity within our basic metaphors. Note that the intrinsic and contingent views merge when we assume that the system's interaction with its observer-controller is negligible, seemingly a reasonable hypothesis in the production setting.

Probably the most well-studied notion of complexity is the one for the computation metaphor. Recall that in this setting a product is represented by a different bit string at each stage of its processing and, as a result, we can regard the entire production system as the execution of a computation; that is, the transformation of one bit string (the initial input) into another (the final product). We then measure the *computational complexity* of the process by the number of distinct operations needed to effect the transformation. In short, a production process is complex if it must do a lot of "work" in producing a given item and simple if the work involved is low. Of course, this view of complexity hinges upon the nature of the computational actors available for carrying out the process. If the actors are each capable of many different types of operations (highly *flexible!*) then the process will be simple; conversely, if each actor can execute only primitive operations (rigid), the process will be complex. This argument suggests an inverse relationship between flexibility and complexity.

If we shift our attention to an (M,R) network as the model of the production system, then our earlier discussion of reestablishable components makes clear the fact that every such network can be thought of as a directed graph. Consequently, any of the numerous concepts of complexity developed in this context can be brought into play [25]. In addition, we also use the fact that every (M,R) system can be given the structure of a sequential machine [15] (although perhaps infinite dimensional). Beginning with the work of Rhodes [24], an extensive algebraic theory of complexity for such objects has been developed over the past decade or so.

In closing the topic of complexity as a system property, let us note the large literature exploring the concept of life, as exemplified in the DNA code, as a language-like system [26]. The complexity of DNA has been considered from several points of view and seeing this concept provide a basis for linking together two of our basic manufacturing paradigms is satisfying. Whether or not the study of system complexity in either context will provide a deeper insight into manufacturing remains to be seen, but it is encouraging to have such a substantive body of past work to draw upon for the task.

D. ADAPTATION AND EVOLUTION

The production system of the future must be not only self-repairing and flexible, it must also be a *learning* system capable of automatically changing its configuration and operating procedures to accommodate shifting environmental conditions. This means that we must have a structure that can *adapt* and *evolve*.

Within the (M,R) system set-up, the machinery for adaptation is already present through the replication function. Basically, we have the equation

$$\text{adaptation} = \text{mutation} + \text{selection},$$

and the mutation process can be represented through a change of the repair map ϕ_f. This, in turn, is effected by means of the replication operation. The missing ingredient is the selection mechanism, which must evaluate various mutations and retain only those that contribute to the system's overall "fitness." For industrial-production processes, this presumably means that we need a procedure for measuring the cost-effectiveness of a given production process operating within a specified environment. No shortage of such measures exists in the industrial engineering/OR literature, so we shall pass over this point in silence.

Contrary to biological assumptions, in the industrial world we assume that "mutations" do not occur randomly, but rather are directed or at least deliberately induced. We have, however, already seen that no type of environmental disturbance (metabolic change) can affect the output of the repair map, so no possibility exists for Lamarckian-type changes within our (M,R) framework. This may or may not be desirable as a feature of production systems and will merit further investigation.

For the linguistic metaphor, the adaptation problem can be given a somewhat different form. Imagine that due to external circumstances (context), a given sentence representing the production of a specific item is no longer "meaningful" and we wish to "correct" the error(s) in the sentence. Thus, we regard the current production scheme as a mutation of the correct scheme we desire and ask how to change the sentence to give it the correct semantic content. When stated in this manner, the problem becomes one of *reversing* mutations to an originally correct sentence. This is exactly the problem faced by designers of self-diagnosing computer compilers that must correct syntactic mistakes in the source language in order to generate admissible target (machine) language statements. As it turns out, this problem can be formally stated in graph-theoretic language, and for general context-sensitive languages its resolution is an NP-complete problem [16]; that is, it involves computing resources that increase exponentially with the size of the problem. Nevertheless, the size of most production processes can be expected to involve sentences that are rather short and, therefore, computationally tractable despite the theoretical barrier of NP-completeness.

Regardless of the specific formal structure chosen to model the production system, certain components must be present in that structure in order to adequately account for the phenomena of learning and adaptation. Because these features are such an important aspect of the production systems of the 1990s, explicitly indicating the ingredients that any adaptive model of production must possess seems worthwhile. Essentially, five components make up an adaptive model [17]:

$S = \{S_1, S_2, S_3, \ldots\}$, the set of attainable structures for the production system;
$T = \{T_1, T_2, T_3, \ldots\}$, the set of operators for modifying production structures. Each $T_i : S \rightarrow \mathcal{S}$, where $\mathcal{S}$ is a set of probability distributions over S;
E = the set of possible external inputs (or disturbances) to the production system;
$\pi : I \times S \rightarrow T$, the adaptive plan that determines what transformation is to be made to the system at time t;
$J : S \rightarrow R$, the payoff for each configuration of the system.

Using the above functions and notation, we have the following adaptive scheme. First, apply π to the current external circumstances and structure to obtain

$$\pi[E(t), S(t)] = T_t \in T.$$

Next, generate $P(t+1) = T_t[S(t)]$. Here $P(t+1)$ is a *particular* distribution over S. Finally, determine $S(t+1)$ by drawing a random sample from S according to the distribution $P(t+1)$.

E. PERFORMANCE

All of the preceding conceptual notions lead up to the fundamental question of how to measure the performance of a manufacturing system. Does flexibility improve performance? Nobody seems to know. What role does reliability play in the overall measure of a manufacturing system's level of performance? Who can say? And so it goes. The plain fact seems to be that no agreed-upon measure will enable us to assign levels of "goodness" for manufacturing systems.

In the spirit of this paper, it would seem reasonable to make a stab at the performance-measure problem by invoking some type of minimality principle measuring performance in a physical or biological setting. Depending upon which of the foregoing metaphors one considers, principles like least time (Fermat), least effort (Le Chatelier), or least energy leap to mind. In a more modern vein, one could also consider principles of least information loss or minimal complexity as candidates for a performance measure. This is clearly a topic requiring considerable further investigation.

The ideas outlined above under points A–D barely begin to scratch the surface of system concepts relevant to manufacturing systems in general and production processes in particular. However, the breadth and richness of even these glimpses into the system structure of industrial processes shows the promise that modern system theory holds for contributing to a deeper understanding of what we mean by a manufacturing system and the ways in which such systems should be structured to accomplish our desired ends. It will take major efforts on the part of many people to carry out the research program hinted at here, but the arguments presented above will, we hope, justify such an expenditure of time and money.

As a final word on the general topic of manufacturing, rather than manufacturing systems, one can raise the overriding question of whether or not the current worldwide flurry of activity devoted to new manufacturing ideas and procedures is an indication of a new growth phase or the beginning of the end of manufacturing as a human enterprise.

Adopting a Spenglerian view, we can consider whether or not manufacturing will follow the path of agriculture and start to display some of the following "downside" characteristics:

1. The government intervenes in order to protect, rescue, and preserve agriculture (manufacturing), running huge deficits to do so;
2. The government establishes extensive R&D programs to advance the technology (e.g., land-grant colleges);
3. Large product surpluses appear, which then have to be "plowed-under" or dumped;
4. As productivity goes up, the number of independent farmers (firms) goes way down—by orders of magnitude; and
5. Agriculture (manufacturing) becomes the very bastion of protectionism.

To see that most of the bad characteristics of agriculture now seem to be present as characteristics of manufacturing in its present phase is sobering. However, it is also heartening to note that agricultural output is also now at its highest level in history, so there appears to be no direct causal relationship between the appearance of the downside features and ultimate systems performance as measured by consumable output. Nevertheless, a general theory of manufacturing should be one that addresses more than just the individual-firm–level questions, containing the seeds of explanation for the emergence and possible modification of these higher-level human and societal issues.

This paper has benefited greatly from the thoughtful reading and carefully argued criticisms of J. Ausubel, R. Ayres, T. Lee, R. Rosen, and E. Schmidt.

References

1. Casti, J. *Dynamical Systems and Their Applications: Linear Theory*. New York: Academic Press (1977).
2. Casti, J. *Nonlinear System Theory*. New York: Academic Press (1985).
3. Kalman, R., Falb, P., and Arbib, M. *Topics in Mathematical System Theory*. New York: McGraw-Hill (1968).
4. Hofstadter, D. *Gödel, Escher, Bach: An Eternal Golden Braid*. New York: Basic Books (1979).
5. Chomsky, N. *Language and Mind*. New York: Harcourt Brace Jovanovich (1972).
6. Thom, R. *Mathematical Models of Morphogenesis*. New York: Wiley (1983).
7. Farmer, D., et al. (eds.). *Cellular Automata*. Amsterdam: North Holland (1984).
8. Baer, R., and Martinez, H. Automata and biology. *Annual Review of Biophysics and Bioengineering* 3:255–291 (1974).
9. Eilenberg, S. *Automata, Languages and Machines*, Vol. A. New York: Academic Press (1974).
10. Wolfram, S. Cellular automata as models of complexity. *Nature* 311:419–424 (1984).
11. Hillis, D. New computer architectures and their relationship to physics or why computer science is no good. *International Journal of Theoretical Physics* 21:255–262 (1982).
12. Manthey, M. Nondeterminism can be causal. *International Journal of Theoretical Physics* 23:929–940 (1984).
13. Chatterjee, A., Cohen, M., Maxwell, W., and Miller, L. Manufacturing flexibility: Models and measurements. ORSA/TIMS Proc., University of Michigan, Ann Arbor (Aug. 1984).
14. Casti, J. On system complexity: Identification, measurement and management. In *Complexity, Language and Life: Mathematical Approaches*, J. Casti and A. Karlqvist, Eds. Heidelberg: Springer (1985).
15. Rosen, R. Abstract biological systems as sequential machines. *Bulletin of Mathematical Biophysics* 26:103–111 (1964).
16. Wagner, R. Formal language error correction. In *Time Warps, String Edits and Macromolecules*, D. Sankoff and J. Kruskal, Eds. Reading, Mass.: Addison-Wesley (1983).
17. Holland, J. *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press (1975).
18. Hatvany, J. Intelligence and cooperation in heterarchic manufacturing systems. *16th CIRP International Seminar on Management Systems*. Tokyo (1984).
19. Gunn, T. The mechanization of design and manufacturing. *Scientific American* 247:87–108 (1982).
20. Goldhar, J., and Burnham, D. Changing concepts of the manufacturing system. *U.S. Leadership in Manufacturing*. U.S. National Academy of Engineers Annual Meeting (1984).
21. Merchant, E. Flexible manufacturing systems: Robotics and computerized automation. *Annals AAPSS* 47:123–135 (1983).
22. Abraham, C., et al. A research agenda for models to plan and schedule manufacturing systems. preprint (1985).
23. Hatvany, J., et al. The computer-controlled manufacturing cell-achievements, possibilities and perspective. *Proceedings of the 9th IFAC World Congress*. Budapest (July 1984).
24. Rhodes, J. *Applications of Automata Theory and Algebra in Complexity Theory*. Berkeley: Mathematics Dept., University of California (1971).
25. Roberts, F. *Discrete Mathematical Models*. Englewood Cliffs, N.J.: Prentice-Hall (1976).
26. Berlinski, D. The language of life. In *Complexity, Language and Life: Mathematical Approaches*. J. Casti and A. Karlqvist, Eds. Heidelberg: Springer (1985).

Received 12 December 1985