Management Services: A Magazine of Planning, Systems, and Controls

Volume 1 | Number 2 Article 7

5-1964

Industrial Dynamics

Bruce Carlson

A new technique for first identifying and then correcting the underlying causes of major trouble spots in an industrial organization has been under development for the past three years at the Sprague Electric Company. Here's a first-hand report on how well it has worked—

INDUSTRIAL DYNAMICS

by Bruce Carlson

Sprague Electric Company

A THERMOSTAT controlling a home furnace . . .
A man driving an automobile . . .
A manufacturing company . . .

At first glance, the three have absolutely nothing in common. But on closer review, they do have one single identifying similarity.

Each represents an information feedback system in which a stimulus—the temperature, another car, a change in orders—causes a reaction. The reaction in turn affects the stimulus. The change in the stimulus then creates a further reaction. The process is one of continual interplay and adjustment, as information flows back and forth within the system.

Each of the three is a closed loop information system, in which one action creates a reaction which modifies the first action. The thermostat, when the temperature drops below a certain level, switches on a furnace. When the furnace brings the temperature up to the desired level, the thermostat turns the furnace off. The man driving down the street automatically reacts when his car deviates from the speed or direction he desires; and his reaction corrects the deviation.

A business follows much the same pattern, except that its reaction time is much slower. A rise in orders will call forth a reaction within the business, but it may take weeks to occur, because so many factors and so many people are involved. The necessary information needed by each of the people to make the correct decisions about the action to be taken has a time lag factor much longer than that present in the case of the man driving down the street. The man can react almost instantaneously; the corporate enterprise will take much longer.

Yet both are dynamic; both do react to stimuli which they, in turn, modify.

This extremely general phenomenon is the underlying basis for a new management theory called Industrial Dynamics. Relatively new, that is. Actually, the theory has its roots in work in electrical engineering which has been evolving since 1937, but we at Sprague Electric were the first to systematically apply it to an industrial situation.

Basically, Industrial Dynamics involves the construction of verbal, graphical, and then mathematical models of the closed loop feedback characteristics of the most important activities of an industrial system. In the models, conditions—or more precisely, delayed and distorted information about conditions—are the bases for decisions that control actions which in turn alter the conditions that are the bases for other decisions, and so on.

The differences between Industrial Dynamics and operations research are the differences between systems engineering and operations research. By a better understanding of the entire relevant system, engineers try to improve its performance. On the other hand, operations researchers traditionally focus on getting mathematically optimum solutions to problems arising in a relatively small part of the system. Industrial Dynamics models are constructed to follow the broad outlines of all important facets of a system and to reveal the causes of behavior that could not be found by examining each of the system's parts separately.

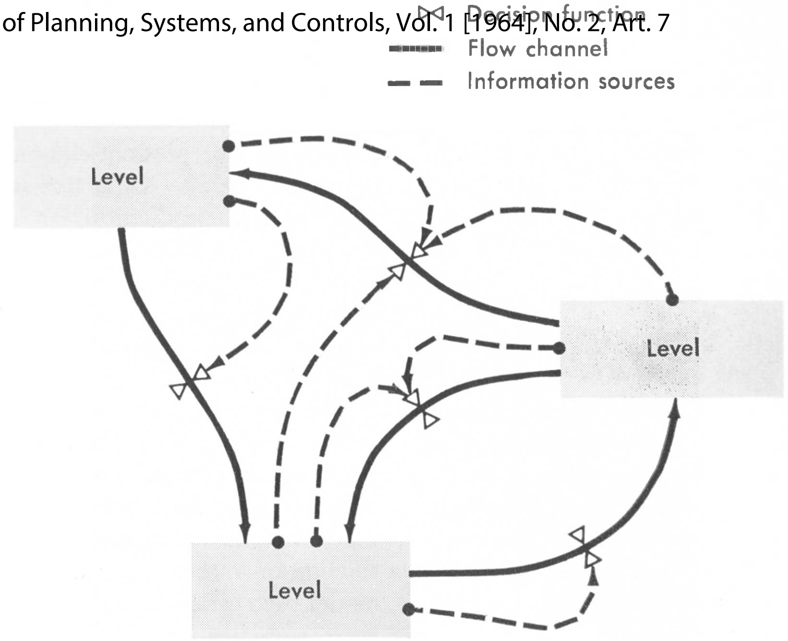

The basic Industrial Dynamics model is based on levels and flow rates. As the contents of one level flow to another level, information about the flow is relayed to decision functions that control the rates of flow.

Goes beyond systems analysis

This may seem like a description of the ordinary systems approach used in business, in which painstaking research into all details of a company's operations eventually results in flow charts that show the company as a total, functioning entity. To a degree this is true, but the Industrial Dynamics approach goes beyond representing the functions within a single company and beyond the flow chart stage to mathematical equations which are simulated on a large-scale electronic computer.

Every ID model has four basic features:

Levels which represent the accumulations at various points in the system at any given point in time. An inventory, a bank balance, the people in a department, are each a level. Looking at them another way, levels exist everywhere there are delays in flow rates.

Flow Rates are the present movements between levels. Flow rates indicate activity; levels measure the state to which the system has been brought by the activity.

An example: an electric power generating station. Its total capacity would be a level; the demands made upon that capacity would be a flow. If demands exceeded capacity, additional capacity would have to be constructed. So flow rates determine levels—as levels do flow rates.

An example: inventory reorders. When stock in inventory (a level) goes below a certain predetermined point, additional stock is ordered from the factory. Movement of goods (a flow rate) from factory (a level) to inventory (a level) will change both factory and inventory levels.

Decision functions or rate equations determine how the information received about levels leads to the decision whether to lower or increase a flow rate. Thus, in the simplest example, an automatic reorder point for inventory would initiate an increase in the flow rate from factory to inventory whenever that point was passed.

Information channels are the media connecting decision functions to levels.

It is the Industrial Dynamics thesis that this basic structure can be used to describe the simple networks that, when all put together, form the company model. Six or fewer networks generally provide a meaningful model of an industrial situation.

The materials network, which represents all flow rates and levels of physical goods.

The orders network, which includes orders for materials, requisitions for new employees, purchase of new plant or office space.

The money network. Here money is used only in the sense of actual cash, with money flow the movement of payments between money levels. The bank balance is a money level under this concept; accounts receivable and price are not included; they are part of the over-all information network which interconnects all the others.

The personnel network, which outlines the company's position in terms of available manpower and utilization of manpower. Obvious levels here would be the labor pool, men in training, men working at the factory. Flow rates would be the rates at which workers were moving from one level to another.

The capital equipment network, which includes factory and storage space, tools and equipment. Flow rates would include the installation of new equipment and production space, and the discard rate of old machines.

Finally, and most important of all, there is the interconnecting *information network.* Obviously none of the five subsidiary networks can exist in a vacuum; decisions in each are influenced by information flowing in from other networks. So the information network is the co-ordinating system for all the others, transferring information about any level to decision points using that information in any network.

An example: A radical change in the orders network will invariably affect the materials network, and could affect the personnel and capital equipment network as well. It should certainly be communicated to the money network. Thus, the information network has the job of tying together the entire company into a cohesive whole able to make a co-ordinated response, just as the nerves in the human body make possible a logical and controlled response to some outside stimulus.

Industrial Dynamics at Sprague

Sprague Electric Company is the country's largest producer of capacitors and other electronic and electrical circuit components. Research and engineering have long been the animating spirit of the company, and profits over the years have been the largest when the percentage of new products has been the highest. This has led to a recognition of the pervasiveness of change, while at the same time management's profit consciousness has resulted in a search for better ways of doing things that carries over into the field of management, as well as product development. In 1957, a joint research project to find ways to improve the operation of one of our important product lines was undertaken by Sprague and the Massachusetts Institute of Technology. Jay Forrester, a member of the M.I.T. faculty and the founder of Industrial Dynamics, was the project leader. The product chosen was a miniature paper capacitor which had been in production for about ten years, and is sold to a variety of customers whose annual usage was relatively stable, but whose total weekly order volume fluctuated very widely.

The first approach taken by Professor Forrester and his team was construction of what would now be considered a conventional mathematical model of the company's production-inventory operations. The model used order-by-order, item-by-item statistical simulation of the Monte Carlo type. Forrester soon became convinced that the approach was not comprehensive enough. The model was not oriented toward the most important managerial decisions concerning the product; the interaction of the company with its market and labor supply. It presented a static picture of a part of what was really a constantly changing situation, and thus could not lead to any significant basic improvement in the system. He therefore persuaded management that a better approach was to construct a model which emphasized the information feedback characteristics of the system, which gave greater recognition to the interacting delays and decisions within the system, and to its behavior as a whole.

The goal was to reduce costs and improve delivery service by stabilizing our inventory and production levels. We are in a business which is not seasonal, but which does have cyclical swings in incoming orders which also appeared, but in greater magnitude, in our manpower and inventory variations. One of the first things revealed by our study of the models was that our old practices actually were a contributing cause of the fluctuations in manpower, inventory and even incoming orders.

How? It was a matter of timing and a human tendency to overcompensate in time of crisis. But these causes were obscured in the multitude of complex day-to-day details of the business. It was intensive investigation and the controlled experiments performed by simulating the models that enabled us to detect them. The roots of the problem were not in our order processing and scheduling procedures, as commonly believed, but rather in our employment and inventory reordering practices.

Delivery cycle vital

For the capacitor and many of our products, the delivery cycle is the vital competitive factor. Because these components have been on the market for a long time, the quality provided by reputable suppliers is about equal. Prices of important producers are also at about the same level. But delivery time is vital. Many of our customers are working on government or industrial contracts where heavy penalties accrue if promised delivery dates are missed. It is essential that they have the components they need at precisely the time they need them.

But there are fluctuations in week-by-week orders, so that we cannot detect a continuing rise in the order flow until several weeks have passed. In our business you can't have a "shape up" every morning, and that is just about what would be necessary if we immediately reacted to every change in order rate. On the other hand, if we react too slowly, order backlogs and consequently delivery delays will put us at a competitive disadvantage.

We were thus faced with a situation that ran roughly like this:

A rising order flow would not change our production level at all to start with. Orders would be filled from inventory as much as possible with the factory manufacturing "to order" components and replenishing inventory as best it could. Over a period of rising sales, however, backlogs would increase as would inventories after a brief initial decline. By the time the over-all rise in orders was detected and steps taken to increase production, average delivery time would have lengthened.

Shortly our customers would become aware of the lengthening, and would increase their lead time for orders so they could be assured of meeting their delivery dates. Customers ordering farther in advance place orders over and above those hey already have in the pipeline, hence causing a total order rise greater than would occur if we could have maintained our normal delivery time. This put an even heavier load on the factory, thus exaggerating what appeared to be the necessary increase in production capacity.

Our inventory policy, which was much like any that is based on traditional rules of thumb—for instance, maintaining four weeks of sales—or, that is derived by conventional OR techniques—as in filling a constant 95 per cent of orders from stock—also aggravated the problem. Why? Because when sales increased not only did actual inventories go down, but desired inventories went up and alert inventory control people placed more than the normal number of inventory replenishment orders on production. So, not only were production people scrambling to keep up with an increase in customer orders, but also with inventory orders.

Increasing production capacity, for Sprague Electric, within broad limits means simply hiring additional workers. However, such workers have to be trained. Thus, in addition to the time it took us to decide to hire additional workers, there was a further lag before we could begin to whittle down delivery delay. During this time customers continued to increase the lead time on their orders, which in turn caused even bigger backlogs, and even greater delivery delays until the new employees and those training them became productive enough to start reducing the backlogs.

At this point there were more people in production than necessary to take care of incoming orders. As backlogs returned to what were considered normal levels, the delivery

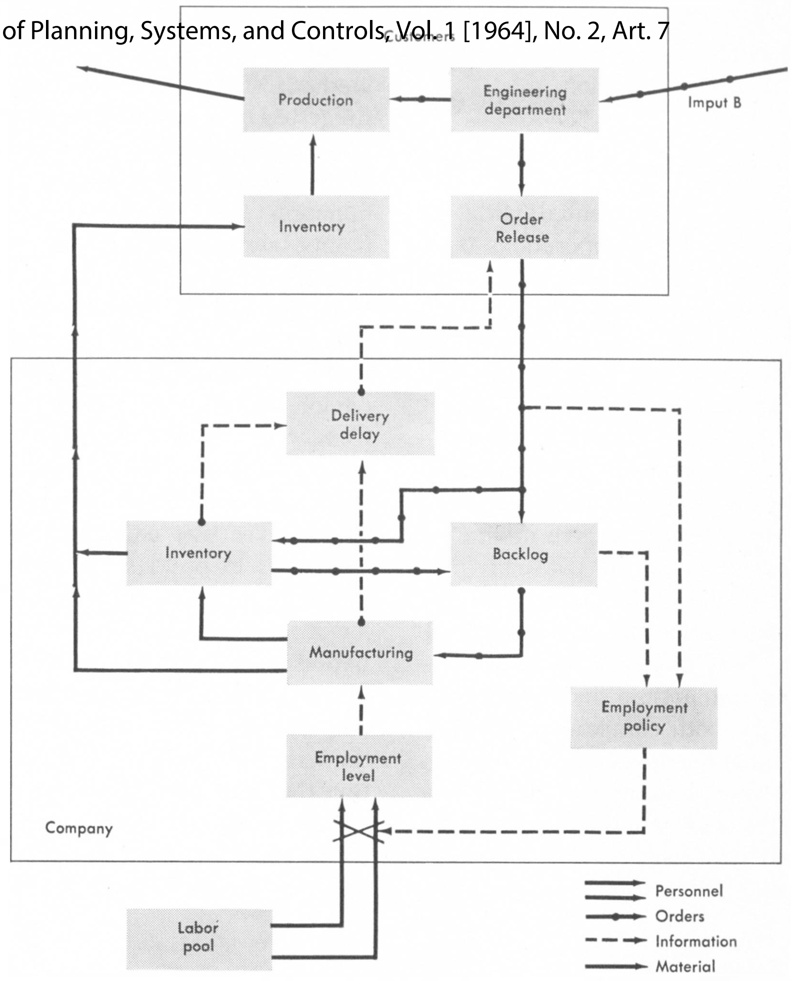

Major system details of the Sprague Electric Company's inventory-production function. Orders flowing into the Customer (Input B, top right) cause the Customer Engineering department to issue orders to its own Production department and to the component supplier (the Company, bottom box). Incoming customer orders are separated into stock orders which go directly to Inventory, and special orders which go into Backlog to be transmitted to Production. Under the old pattern, as Backlog of orders built up, it was reflected in Employment policy, which caused the Employment level to rise as the Company drew on the Labor pool by raising the rate of hiring and training more workers for Manufacturing.

delay shortened and, consequently, after a time customers began ordering less far ahead.

There was, in other words, a vicious cycle, producing oscillations in production, inventory, and manpower far greater than the total change in sales. In our effort to take corrective action, we were actually aggravating it. In their effort to take corrective action, our customers, too, were further aggravating the problem.

This showed up very clearly in the models. We found that our inventory was peaking shortly after production peaked, and was rising while production was rising. Conversely, inventory was falling while production was falling. This, of course, amplified the production downswing, so that production went lower than sales. In other words, our inventory policy both increased the peaks and decreased the valleys of our production fluctuations.

It was postulated that a system which was more sensitive to incoming orders, less sensitive to backlogs, and in which inventory orders were increased when sales were falling and decreased when sales were rising would help dampen the extreme fluctuations in production. This proved true when tested with the models. The result: a greater stability in inventory and employment, and, concomitantly, a better delivery delay situation.

Implementation

We first experimented with the program in actual operations on a tentative basis. Implementation of the new policies began in early 1961, and the weekly employment decision rule was being used on a routine basis by September of that year. Although middle line management had been fully informed about the project from the start, and was often asked to give advice on various points, it did not really participate fully in the project until the implementation phase was reached. In the early phases of the implementation, much time and effort were spent by both company and M.I.T. personnel in educating line management by explaining in detail the new policies and what they were designed to achieve. No one was forced to follow the new procedures; each executive involved was personally sold the idea of the new methods. In return, the managers suggested minor modifications to the policies and pointed out a number of additional problems—the determination of inventory reorder points, systems changes necessary to generate more easily the information required to operate the new policies, and so forth—the solutions to which improved the systems behavior. Actual results of deviations from the new methods were checked against results predicted by the model under the new policies. The answer was conclusive. If we continued to build some inventory during a period of falling sales, and let it temporarily drop during a period of rising sales, we would be able to stabilize both our work force and our inventory, and cut down the extreme oscillations in production that had been characteristic of the product line. Over a period of time, considerable interest developed in the application of the new employment decision rule and it was followed closely.

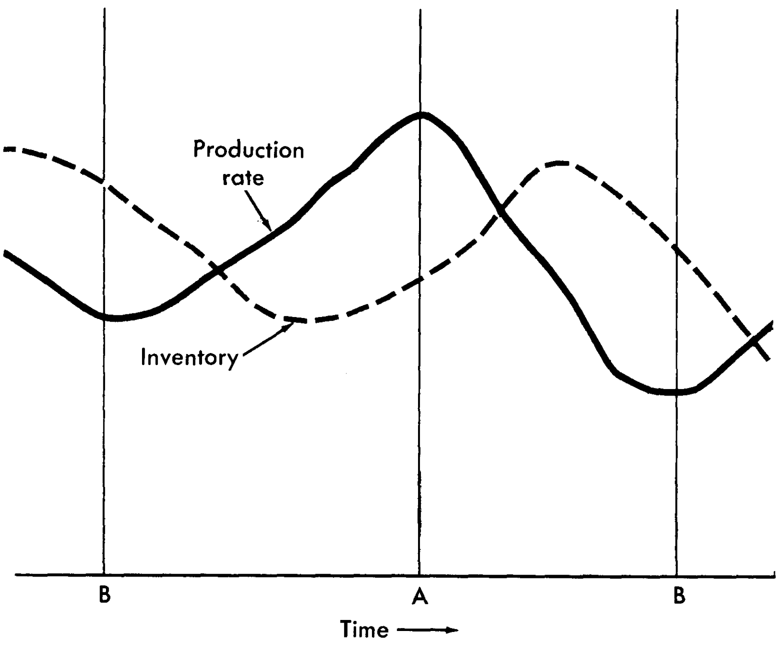

Figure illustrates production-inventory pattern under old system at Sprague. As production rose to meet orders, inventory levels followed rise. But the increase in production and inventory automatically resulted in drop in orders as backlog was reduced. Thus production peaked as orders were declining, and inventory reached highest point after production had begun to fall off.

The production scheduler was given a priority list of program stock items that could be produced whenever factory capacity exceeded that necessary for customer orders. On the other hand, customer orders are always given priority over inventory replenishment orders. Finally, we changed our policies so that total authorized employment each week was based on long-term average sales and inventory adjustment over a long period of time (to better balance the absorption of sales fluctuation between inventory and employment).

We should make clear at this point that the new policies, and particularly the employment decision rule, did not replace the existing system completely, but rather served to provide an additional input to decisions that were then and still are being made by operating management. In fact, there have been and still are occasions when the rule is disregarded, because it is very difficult for line managers conditioned by years of experience to reacting intuitively to sudden changes in business conditions to accept with complete faith the relatively slow reaction times called for by the policies of the I.D. model. The important thing, however, is that despite occasional deliberate deviations, the new policies continue to influence the decision maker and he to influence the policies, which have undergone a number of minor modifications.

Results

On the basis of our experience to date, we can report a number of tangible accomplishments which we believe are directly attributable to the application of the model. These are as follows:

1. Comparison of data before and after implementation of the new policies shows that productivity, measured in units per man-hour, increased by 12 per cent. This we believe is partly the result of the more stable employment called for by the

new decision rule, and partly the result of aggressive methods improvement programs.

2. The production cycle has been slightly shortened, as a result of the higher productivity and improved scheduling procedures.

3. The inventory level is being used to absorb factory pressure in periods of peak sales by cutting back inventory production and allowing the inventory level to fall. In one period of peak employment before the changes, inventory was rising at the rate of 10 per cent of the average sales rate. At peak employment after the changes, inventory was falling at the rate of over 11 per cent of the average sales rate, and this contributed to a variation in employment that was considerably less than would probably have occurred under the old policies.

4. Inventory is better balanced, because of certain procedural changes suggested in the course of the study and the use of computer-established reorder points.

No attempt has been made to measure exactly the effect on profits of the new policies. This would be a very difficult and costly undertaking requiring extensive changes in company-wide accounting systems, which are not deemed worthwhile in view of the fact that the study has been confined to one product line among approximately fifty on which profit information must be compiled. Instead, all concerned have agreed that the record of operations since the new policies were introduced clearly shows improvements. Some of these improvements have, we believe, resulted from the application of more conventional systems and procedures and quantitative decision-making techniques, but it is also probably a fact that these changes stemmed in large measure from the insights provided by the model simulations.

In one major respect, however, the Sprague Industrial Dynamics project has not had the predicted effect. This was in the area of long-term fluctuations in recurring orders and the interaction between Sprague's own actions and those of its customers in a closed-loop feedback system. It had been assumed from the start of the project that the long-term fluctuations arose in large part because the customer tends to follow a policy of ordering farther ahead as the company's service delay becomes long, and of holding orders back when the delay becomes short. In other words, the company's employment decisions which directly affect the delay time are reflected in the customer's ordering rate. Thus, by modifying its own policies along the lines of the model, it was hoped to damp out fluctuations in the incoming order rate. We are sorry to say that this has not been achieved; after more than two years, incoming orders for the product in question are fluctuating as widely as ever, and it has been necessary to modify the employment decision rule in recognition of this fact. The importance of the feedback concept in virtually all socio-economic systems is one of the cornerstones of the Industrial Dynamics approach. Although this is probably a valid concept, our project does not, to date, support the part of the philosophy which implies that a relatively few easily discernible factors interact to form feedback loops which persistently dominate the behavior of a given system.

Evaluation

Let us attempt an evaluation of Industrial Dynamics in light of the Sprague experience and touch briefly on some of the problems we see. The three most important ones, in my opinion, are (1) the scope of Industrial Dynamics models; (2) the level of aggregation of model variables; and (3) validation of Industrial Dynamics models.

With respect to model scope, we are not sure that we have learned how to determine what the scope should be, or whether the amount of detail that was finally included in our model is the right amount. In this stage of the development of Industrial Dynamics, primary reliance must be placed on intuitive judg-

Industrial Dynamics, although a relatively new technique, has already created fierce controversy. Its backers feel that it is far more significant than the more conventional operations research methods, because it suggests solutions to more basic problems than can be attacked by O.R. techniques. The difference between the two approaches is the difference that exists between strategy and tactics, in this view. Many O.R. experts, on the other hand, do not feel Industrial Dynamics has proved itself in any way . . .

ment in determining scope. For example, because the Sprague model omitted explicit representation of the company's competitors, there is some question whether the company's market is adequately represented.

Another problem arises from the necessity to include, in order to have a closed feedback loop, a model sector representing the company's customers, about whose operating policies little is actually known. As we pointed out earlier, we have postulated certain apparently logical responses by the customer to the delivery delay he sees on Sprague's part, but we do not know enough about what determines his ordering decisions to say that our customer sector is a correct representation of what really happens. The fact that operation of the new policies has not yet had any noticeable effect on the fluctuations in incoming orders reinforces this feeling, although it may be that insufficient time has elapsed for the feedback effect to be noticed.

With respect to the aggregation of variables, no systematic way has been developed for aggregating nonlinear elements of the system. For example, in the Sprague model we are aggregating orders from different kinds of customers in a number of different industries, and we are aggregating orders calling for immediate delivery and orders calling for extended future deliveries. We hope this is justified because the differences among these various categories are not dynamically important, but there is no way of being sure of this. With all the random fluctuations in the available information about real-world phenomena, it is virtually impossible to detect underlying causal mechanisms from existing data. We are therefore forced to rely on the opinions of operating people who are close enough to have formed opinions about what is important in a given activity, but who are also close enough to it to have biased views.

With respect to the validation of Industrial Dynamics models, Jay Forrester has stated that:

"The significance of a model depends upon how well it serves its purpose. The purpose of Industrial Dynamics models is to aid in designing better management systems. The final test of satisfying this purpose must await the evaluation of the better management. In the meantime the significance of models should be judged by the importance of the objectives to which they are addressed and their ability to predict the results of system design changes."*

We believe the present Sprague model has "served its purpose," in that it has led to the design of improved management policies relating to important operating objectives of our company. It has demonstrated an ability to predict the results of system design changes in some areas, but has not succeeded in damping out the large fluctuations in the order rate. However, we are sufficiently satisfied with the results to date that we have extended the use of policies similar to those tested by the model to at least one other product line, and probably will extend them to still others in the future.

We have the impression that Industrial Dynamics as a discipline is increasingly concerned with the more intangible aspects of industrial behavior. This may be because one is more immediately aware, in production-distribution systems of the type represented by this first Sprague model, of problems of scope, aggregation, and others I have not mentioned. Moreover, in such systems, available data are apt to be distorted, making it difficult to sort out the really significant causal relationships that must be represented in a model. Another difficulty in applying Industrial Dynamics to these more tangible types of problem is that the system changes suggested by Industrial Dynamics models may be in direct conflict with other more conventional control techniques. It may, therefore, be very difficult to apply the new policies in one area of a company's operation without applying them throughout, which may well entail dropping some cherished control procedures and devising new ones to more adequately measure system response.

The Industrial Dynamics project at Sprague Electric has followed a similar evolution. In a general way, we have applied Industrial Dynamics to higher level problems in contrast to the relatively low-level production-distribution system represented by the model discussed here. Higher level problems deal with such things as new product growth, divisional growth, or total

*Forrester, Jay W., *Industrial Dynamics* (M.I.T. Press and John Wiley & Sons Inc., 1961), p. 115.

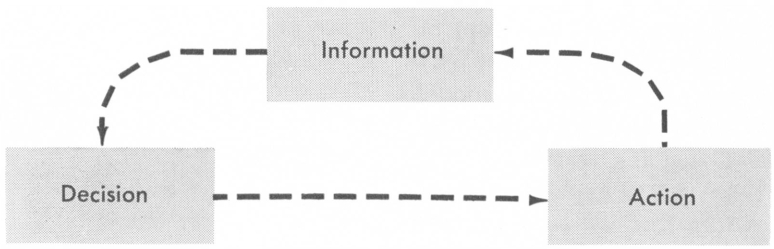

Diagram illustrates decision and information feedback pattern. Information influences decisions which determine actions. Results of actions taken become information which in turn influences subsequent decisions. These in turn affect subsequent actions in a continuous chain.

corporate growth. In general, they are more important, more long range, and more difficult to solve. Most significantly, they deal more with intangibles, and the benefits of a study are likely to come more from insights gained in making a study than from any mathematical decision rules that may emerge.

When applied to solving tangible problems, we believe Industrial Dynamics is slower and more costly than conventional operations research techniques, and to date has probably yielded smaller payoffs, although it has certainly pointed out some important shortcomings of certain widely accepted analytical procedures. One thing we are certain of is that Industrial Dynamics will profoundly influence the development of other management science techniques. This has happened already at Sprague Electric, where, it should be added, we are continuing to devote a portion of our management research effort to more conventional areas of management science.

It appears to us that the conceptual relevance of Industrial Dynamics is being more widely accepted today in the business schools than it is in industry, and we hope that the experience of Sprague Electric Company and others with this powerful new philosophy will help to bring it to the attention of more and more business managers.

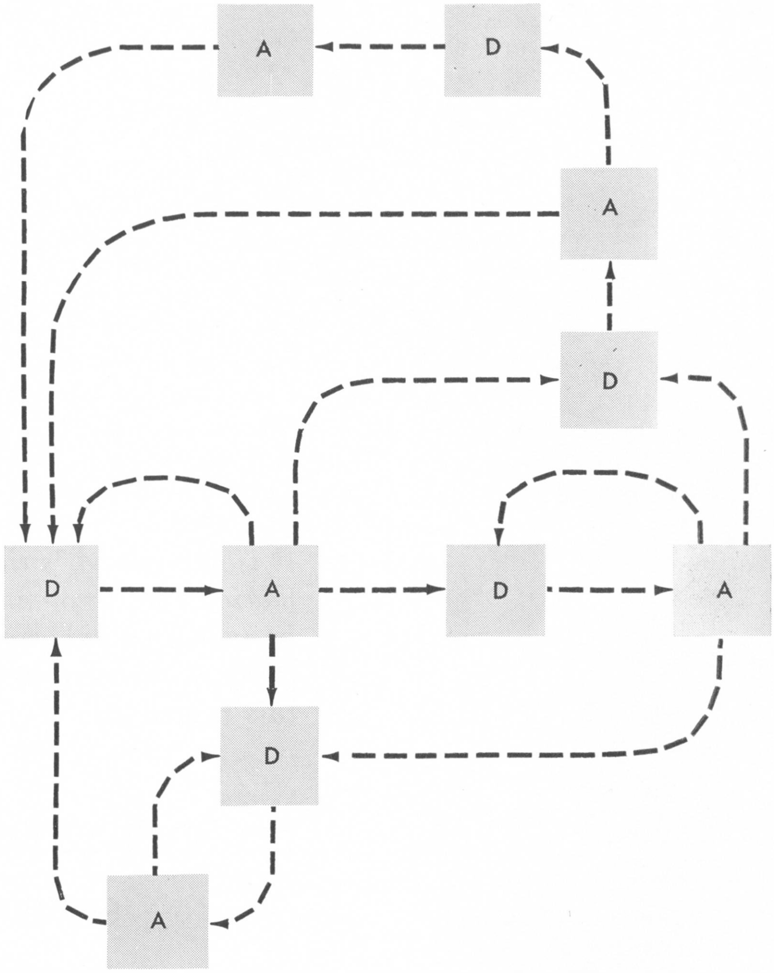

In any industrial system, the simple information-feedback loop shown on facing page becomes an infinitely more complex multiple loop, interconnected system. Decisions are made at many points throughout the system, each of which results in actions that create information to be used at various other decision points in the complex.

One of the major advantages claimed for Industrial Dynamics by its partisans is that it is far more easily understood by management all down the line than some of the esoteric mathematical techniques employed in conventional operations research applications. All equations in Industrial Dynamics, for example, are expressed in a form of shorthand representing common business words. Thus: one term of an equation might be IFAC, which would stand for Inventory, Finished, Actual at Customer (Units). Since many equations encompass time, three arbitrary symbols, J, K, and L are used to represent periods in time. J represents the previous solution time of an equation, K represents the present moment, and L represents the next solution of the equation. (The time interval between J and K and between K and L is always fixed, and is known as Delta Time.)

Thus, a complete equation to determine the state of a present level would be:

IFAC.K = IFAC.J. (DT)(UAIC. JK–URMC.JK)

Meaning of the terms:

IFAC.K—Inventory, Finished, Actual at Customer (Units) at present time, K

IFAC.J—Inventory, Finished, Actual at Customer (Units) at past time, J

DT—Delta Time (weeks)

UAIC.JK—Units Arriving at Inventory at Customer (Units/week) during the time interval, JK

URMC.JK—Units Received in Manufacturing at Customer (Units/week) during the time interval, JK

In other words, inventory at customer at present time is equal to inventory at previous solution time plus the difference between additions to inventory and shipments from inventory during the interval from previous solution time to present moment.