Proceedings of the 1987 Winter Simulation Conference
A. Thesen, H. Grant, W. David Kelton (eds.)

STRUCTURED FEEDBACK: AN APPROACH TO DEVELOPING KNOWLEDGE OF A SYSTEM DURING INITIAL MODEL DEVELOPMENT

Patrick J. Starr
Mechanical Engineering Department
University of Minnesota
Minneapolis, Minnesota 55455
612-625-2315

Thomas M. Keenan
Applications Engineer
Control Data Corp.
Bloomington, Minnesota
612-625-6175

ABSTRACT

This work describes a proven approach to the task of initial model development for a manufacturing simulation. The approach recognizes that model development engages two groups of people in a mutual learning experience: the model builders, whose initial task is to develop an understanding of the system, and the manufacturing personnel whose initial task is to communicate knowledge of the system. The approach utilizes the fact that a model can be expressed in many forms: narrative, graphical, mathematical, computer code, and user interface. In much of the literature, the term "model" is synonymous with the coded form. However, the narrative and the graphical forms constitute excellent means of facilitating the mutual learning experience. The narrative-graphical model is developed by an iterative means called "Structured Feedback", wherein the modelers develop a sequence of models, each being subjected to critical review in meetings with manufacturing personnel. The result of the Structured Feedback approach is a system description that all parties agree to and understand. The approach is described in the context of an electronic workcell study. Elements from the sequence of narrative-graphical models are shown. Guidelines for narrative and graphics development are presented.

1. INTRODUCTION

Model building requires many decisions to be made. Some of these include the choice of objectives; the boundary; elements and interactions; level of detail; mathematical representation; solution technique; validation arguments and experimental design. These decisions are influenced by the knowledge of the system which is under investigation. As one's knowledge grows, so does one's ability to recognize finer distinctions that influence these decisions. That is, the system boundary may change or the level of detail necessary to describe an operation may increase as the modeler develops a fuller understanding of the situation. In this work, we address the process by which the model builder develops knowledge of the system. It is a process that has been performed by anyone who has constructed a model in an applied environment, yet it is rarely mentioned in the literature. It is acknowledged by various authors (e.g., Shannon, 197), but there is little information about how to embark upon; record; modify; and communicate one's understanding of a system during model development. Here, we take a close look at the process and describe a format to express the knowledge along with the means to update and communicate one's understanding.

This process will be discussed in a specific situation, with the amorphous objective of identifying ways in which simulation could assist in planning and operations decisions for an electronic assembly workcell. It provided the opportunity to explore a wide range of possibilities. Thus, even though the remarks are based on an example system, they have more general applicability. We assume a manufacturing setting, in which the modelers have their foot in the door. We assume that an explicit objective has not been stated yet: developing clearly stated objectives is part of the problem. However, we do presume that a simulation model is an intended outcome.

2. THE ACTORS AND TASKS

Model development will engage two groups of people: manufacturing personnel and the modelers, and each has a different viewpoint. The modeler is a generalizer, and sees systems, processes, flows, broad boundaries, freely flowing information, many alternatives and is not initially interested in numerical details. The client-user is specific, seeing these parts, these machines, this staff and this problem; information is restricted; boundaries and decisions under one's control are equivalent; numerical data is hard to get and is meaningful; many "alternatives" are unthinkable. Furthermore, each has a different view of the capabilities of a model: the modeler sees it as a tool to explore alternatives. The client may be skeptical: he does not know what will be in the model; how it works; or how it could possibly help his situation. Model development should engage these groups in different ways. The modelers should:

1. develop an understanding of the system being investigated,
2. communicate realistic capabilites of simulation,
3. describe the resources needed to develop a simulation,
4. develop realistic objectives by proposing "what-if" situations.

The clients should:

1. communicate knowledge of the system,
2. learn the capabilities of simulation,

3. develop realistic objectives by suggesting possible "what-if" situations.

The task obviously is for these groups to get together to engage in a mutual learning experience. The following offers both a format and a process to facilitate this learning.

3. NARRATIVE-GRAPHICAL MODEL FORMS

In much simulation literature, the term "model" is taken to be synonomous with "computer code". It is well to recall that a model can be expressed in many forms, such as:

1. Narrative ("words")
2. Graphical ("boxes and arrows")
3. Mathematical ("symbols and logic")
4. Computer Code ("a language")
5. User Interface ("menus, animation")

Each form has its own strengths and weaknesses. Obviously, for simulation, the coded form is essential, since it "runs" and provides the time responses and subsequent statistics on the variables of interest. However, few people (clients) understand the coded form of a model. It is the narrative and graphical forms that provide a mutually understood system as a basis for communication between modeler and client. Since the clients know the system (although any one person may only know a piece of it) and the modeler is to learn it, it is the modeler's responsibility to select model forms that will best communicate his understanding of the system back to the clients.

The graphical representation of manufacturing systems is nothing new, but in modeling it can play a crucial role in communicating system elements and structure. It also assists in explaining what happens in a simulation. Discrete parts are followed through the various processes, and are either waiting for processing or being processed. The waiting times and processing times accummulate to yield throughput time and perhaps "lateness". Their effect on the utilization at each station is noted. The procedure is repeated for many parts; statistics are generated over the part population and over time. In this way, the flow in a station-route diagram replicates the generalized process of simulation. A graphical model shows <u>what</u> is in a system and the narrative describes <u>how</u> it works.

There are many ways to organize a narrative; Figure 1 shows the outline that we found useful. In addition to the items listed, the narrative is interspersed with paragraphs that describe appropriate Data Needs and What-If Situations. These are included in the narrative rather than in separate sections, because they require a certain level of detail of system description, a context, to be meaningful.
A later section will show examples from a narrative, but first we will address the means of initiating and altering the narrative. A one-shot attempt at understanding the system will not work. The modelers and clients must be engaged in an <u>ongoing process</u> in which the communication becomes more efficient; the objectives and data needs become clearer.

<u>STRUCTURAL INFORMATION</u>
(What is there? What is planned?)
- .Boundaries of flow, control and time.
- .Station definition and categories.
- .Products, parts, assemblies: number and categories.
- .Routes.

<u>OPERATIONAL INFORMATION</u>
(How does it work? Narrate the normal flow.)
- .Station operation.
- .Loading patterns, batch sizes.
- .Staffing.
- .Processing times.
- .Capacities and blocking.

<u>PROBLEM SITUATIONS</u>
(Flow stoppages and possible causes)
- .Bumpy loading.
- .Part shortages.
- .Down equipment.
- .Capacity increases or shifting.
- .Quality changes.

Figure 1: Outline of Narrative Model

4. STRUCTURED FEEDBACK

Structured feedback is a process of interaction between the modelers and clients that begins with a meeting in which the modelers describe:

1. modeling and simulation in manufacturing;
2. model building decisions; and
3. elements of a graphical-narrative model.

Following this meeting, the modelers interview the manufacturing personnel and examine written materials, flow charts, etc. The modelers apply their "art" and produce the first graphical-narrative model, M1, following the format of figure 1. It is circulated to the manufacturing personnel and a feedback meeting is held during which the modelers describe M1 and the manufacturing personnel critique the model. Deficiencies of M1 are noted and the model is revised as M2. The revised model is critiqued in a second feedback meeting, and the process continues, producing a sequence of graphical-narrative models, until one is agreed upon to describe the system and address the "what if" situations of interest.

Some benefits of this approach for both client and modeler are:
-All participants develop a common understanding of the system under study.
-A variety of "what-if" situations are described, thereby fostering a clear statement of model objectives. (The "objective" is to answer <u>these</u> "what-if" questions).
-Data needs are clarified as well as the role the data will play.
-One form of documentation is produced.

Additional and less obvious benefits of the structured feedback approach to a a graphical-narrative model follow:

The graphical-narrative description is easily translated to pseudo-code in that the flow of products through stations parallels the program structure. The modeler can translate nouns from the narrative into entities, attributes and resources; verbs become operations and decisions.

During the process of describing, graphically diagramming and narrating a system, inevitably the system's boundaries will be broached. The boundaries may be departmental geography, jurisdiction, personnel or time frame. For instance, the process narrative may assume the process begins with the arrival of kits on the floor. It may not extend to the jurisdiction of Materials Management, which oversees the issue of kits. A narrative description of the system can foster better communication about such overlapping areas of concern. Resolving boundary issues gives rise to what-if situations, as we will see in further examination of these issues in the workcell study that follows. Boundary issues can broaden the client's scope of understand of the whole system.
The sequence of graphical-narrative models is a record of the evolving understanding of the system. It provides a record that remains in-house after the project ends. It is a basis for model enhancement as well as for further modeling projects. The knowledge acquired doesn't walk out the door with the modeler.

This approach contains a means of model verification. The narrative is an agreement between client and modeler on the three basic elements of the system (Fig. 1). In participating in its creation, the client has presented shared experiences and postulated problems which are incorporated into the model; and he/she has agreed to their representation as accurate. The client has been able to bring his anecdotes and concerns to the model directly without relying on the accuracy of interpretive coding.

The modeler becomes proficient in understanding the entire system, sometimes more than any individual within the system. This is a distinct advantage to the modeler in his ability to serve the same or other clients in the future.

5. APPLICATION OF THE STRUCTURED FEEDBACK PROCESS

The following describes an application of the structured feedback process to model development in an electronic assembly workcell. Selected elements of the graphical-narrative model will be described and in some cases the narratives will be quoted. The workcell was chosen by the client as a suitable environment to explore the use of simulation to assist in design and operations. The workcell was in the process of redesign; products were in the prototype stage. The client's intention was to finish the development of the production process in tandem with the design of the product itself.

As outlined in principle above, the process began with a workshop to introduce the project. Simulation in manufacturing and the structured feedback process were introduced to the client. Interviews of many manufacturing personnel followed in which we recorded the following findings:

- Dated process flow diagrams were available for each product.
- Alternate workstations were being considered.
- Set-ups were needed at some stations.
- A robot was being scheduled for use.
- Some cell operations were shared with other products.
- Projected requirements increased with the implication that maximum production capacity <u>must</u> be increased.
- Part shortages were a concern.
- Rework operations occured for most boards.
- Uneven loading of a cell was considered a problem.
- There was a desire to move to a JIT operation.
- Throughput time was a key measure of system performance.
- Excessive work-in-process (WIP) was undesireable.

Many of the items listed above could be stated as what-if situations and a simulation could explore their effects upon throughput time and WIP. We were encouraged and proceeded with the structured feedback process.

Three narrative-graphical models were produced, documented and presented at feedback meetings for critique. The models will be referred to as M1, M2 and M3. The following will include excerpts from M3 with parenthetical remarks about the differences between that model and M1 and M2.

Structural Information

The process flow diagrams for each product formed the basis of the structural information. They show the series of operations that a product experiences, but do not convey the difficulty of implementing such operations on the floor. We created a station/route diagram, Figure 2, for a typical circuit board. The diagram was the graphical model; it underwent various changes during the process. It was well-received and to our surprise it was new: this branch of the company did not use such diagrams. The diagram became a useful tool for discussing system boundaries; where the flow came from and where it went when it left the page. Initially, all PF operations were included, but for M2 operations were restricted to only those in the workcell. In figure 2, most stations have two numbers: the one in the circle is a station identifier while the other one indicates the number of different process flow operations that occur at the station. Stations without an identifier are for parts preparation. Visits to the cleaning station are indicated by the "loops" attached to each station where cleaning follows.

Figure 2 appears to be just another textbook example, but much judgment was entailed in its formulation. Stations represent physical locations where operations are carried out. Some stations were associated with a single operation; some had multiple visits for the same operation; and some represented a group of operations being performed by one or more operators in close proximity. In these latter cases, the definition of a station became more conceptual. In M1 we grouped what was believed to be similar operations. In M2 we found that certain technology was being considered that would perform selected operations, so they should not be grouped. Later we would find that for certain groups of operations, a worker will move among a set of specialized benches, so the worker-bench combination became a station. An important point, however, is that the graphical representation can accommodate all these distinctions, as boxes and arrows are ambiguous. The narrative is required to fill in the detail. Thus, one does not need to represent a myriad of exceptions or details in the graphic illustration.

The description of products as structural information is merely identifying their categories and routings. Since each had its own PF diagram, the needed information was available for M1; Figure 2 shows a typical route. For M2 it was learned that for products comprised of two circuit boards bonded to a heat sink, an equal supply of each board would not be available for issue at the same time. This suggested the product should be a circuit board rather than an assembly; it also raised operational issues.

Operational Information

This portion of the narrative described the movement of a typical circuit board among the stations and specified the data needed and possible what-if situations.

Station operation in M1 was treated as a multiple server queue system with no communication to the rest of the system. The server or technology sits and waits for work; it has no knowledge of where the work came from or where it is going. The station is utilized to varying degrees and idle cap-

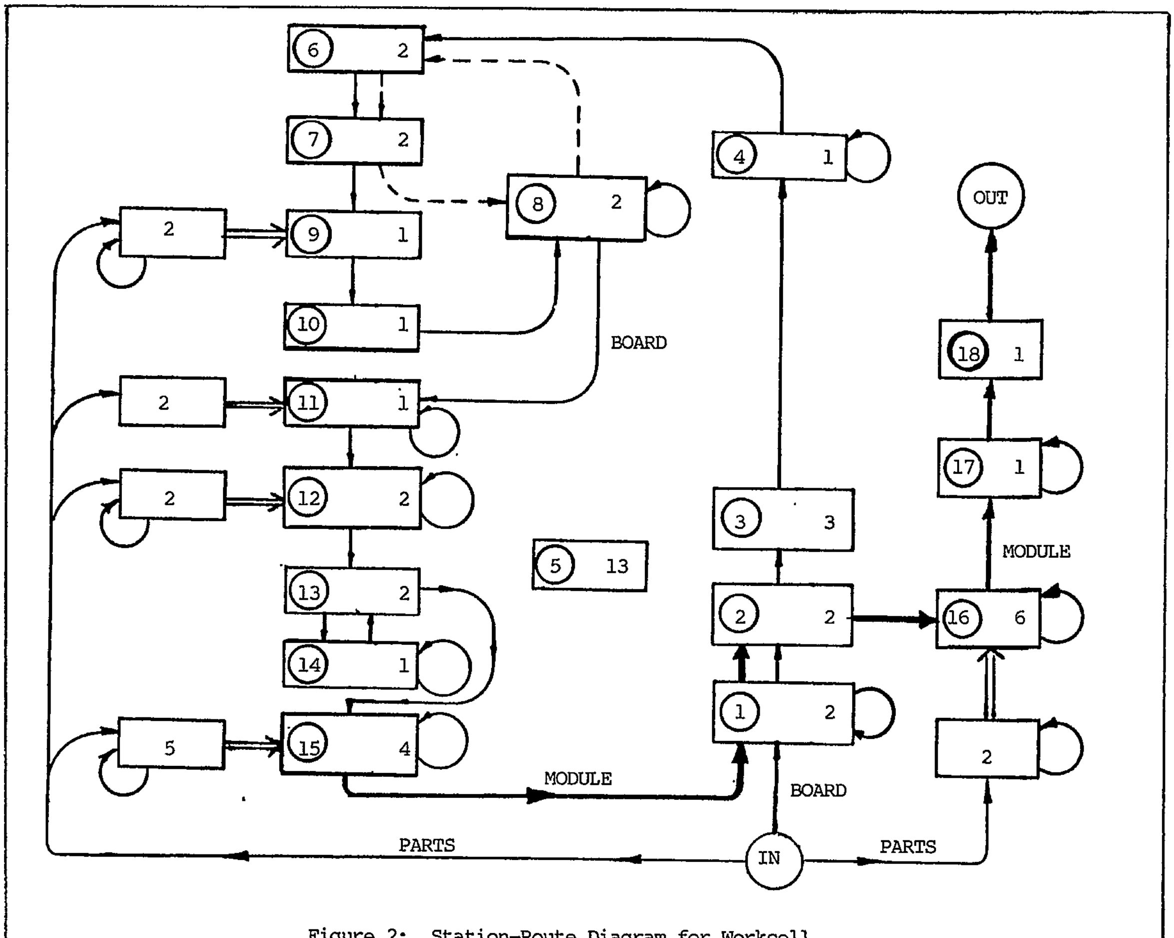

Figure 2: Station-Route Diagram for Workcell

acity is wasted. This approach, when spelled out, can serve as a baseline to be expanded as necessary. In M2 and M3 we found that worker groups serve groups of stations and can move about as needed. In fact, the need for cross-training of workers was a topic of concern. Also, there were cleaning and soldering stations to which a worker brought a board and waited a brief period for self-service. In such cases, servers of one operation became customers of another. Also, at assembly stations where parts or boards are joined, an insufficient part mix results in either a stoppage or processing another (lower priority) board. Materials handling between stations between stations was accomplished by the operator of one station hand-carrying boards to the next station.

Cell operation is intimately related to product type and gives rise to many what-if situations. The following is a portion of the M3 narrative that describes the interaction of product type with cell operation. This description was absent from M1 and evolved during the feedback sessions and subsequent interviews.

Scheduled production includes five different boards from which three different modules are assembled. One of the modules consists of a single board bonded to a heat sink; the other modules are comprised of two boards which sandwich a heat sink. The individual boards of the two-board modules are referred to generically as the "A-side" and "B-side". The transformation from boards to module occurs at station 15, where the heat sink is applied. The presence of the two-board module implies that there may be disparate loading on either side of station 15. That is, if the appropriate A and B sides are not available, station 15 and beyond may have little or no load; while up to station 15 the loading may be relatively uniform. This situation suggests some relevant issues, since it will dramatically affect cell throughput and WIP. If scheduling is done sequentially, with many lots of A sides followed by B sides, stations 15 onward may operate in a "boom-bust" mode. The simulation could identify the degree to which this occurs and be able to explore the influence of board mix.

Partial Builds refers to initiating construction of a board for which the kit is short of selected parts. The board follows its usual route until the inspection operation station 13, or it may pass this station in spite of short parts (a policy decision) and continue on as a module until station 17. One or more supplemental stations could be inserted that would (1) install the needed parts; (2) inspect; and (3) rework.

The situation is easily modeled in terms of stations and routes. The most straightforward would be to define five additional boards as "short part boards". These boards would have two route options:

Option 1: Hold boards with short parts. This would introduce a delay prior to station 1 until parts arrive. When parts arrive, an impulse load of jobs would occur.

Option 2: Initiate work on boards with short parts. This would "buy time" and load the floor. Some process times would be reduced. The inspect and rework at stations 13 and 14 would have to pass the board without all the parts. The module would be held at the supplemental stations after station 16 (or in some storage area) until all parts arrived. Once the parts arrived, another option is faced: install parts at the supplemental stations or reroute back to the earlier appropriate station. Installation at the supplemental station may take a very long time, while rerouting may overload a previous station. In either case, a "pulse" load will occur when the parts arrive. The simulation could describe these options for short part operation.

Prototype Boards utilize the operations associated with stations 6-10. These can be modeled as additional boards that are only routed through those stations. What-if situations would entail scheduling of these prototype boards and their effect on the throughput of regularly scheduled production.

The preceding paragraphs, taken from the M3 narrative, describe operational "what-if" situations associated with the different products. Note that no explicit questions are present. Rather, the content conveys the modeler's understanding of the situations and includes remarks regarding the capability of simulation.

Cell operation is also influenced by the loading. Model M1 did not include the arrival of parts into its system. Boards and parts simply "arrived". Floor managers raised the issue of "bumpy" loading, uneven arrival of parts into the system. We then interviewed Materials Management personnel and examined the actual and the theoretical situations that occurred on the floor during the kitting and issueing process. We included in the narrative a description of the build schedule and the pertinent events that occur in the three months prior to releasing the kit on the floor. Out of this process, we were able to identify two general strategies which were stated in M3 as follows:

a. Take the projected kit issues and staffing as "given" and explore management and techical options within the cell to attain (or reduce) the estimated build time.

b. Presume the issues requirements and staffing estimates have allowable ranges, and use the model to find "good" loading and staffing patterns that fall within the ranges. A temporal dimension could be employed, in which the schedule of completed boards and staffing could be viewed as monthly goals, and the model could be used to identify weekly loading patterns to attain the monthly goals.

These strategies differ in their view of the system boundary. In the first one, no communication with Materials Management is needed; workcell management strives to do the best it can under specified loading. In the second, the targets may be altered, but the monthly goals are preserved. This presumes that Materials Management would be amenable to changing the way weekly loads are established if the model demonstrated some advantage in doing so. This is a reasonable view of the role of a model, but it may be an unreasonable view of the boundary between Materials Management and Workcell Management.

Data Requirements for simulation fall into two categories: one is the bare minimum that is needed to describe a base level of operation, and another is needed to implement various what-if situations. Thus, data needs are intimately related to the objectives of the study. Also, it is important to point out to the client that, for most parameters, one merely needs "reasonable ranges", since simulation is an experimental approach and is easily altered.

We sought the information needed to simulate a base level of operation, following model M1. Available data included the proposed monthly loads for two years; estimates of processing times (but no indication of variability); capacity estimates for technologically limited stations; and current batch sizes. As the narrative expanded through M2 and M3, more what-if situations were incorporated into the narrative in the form of logical statements such as "when these parts are short, but have been on back order for fifteen weeks, issue the load". Figure 3 summarizes the parameters and logical statements needed for the narrative model. Those marked with an asterisk are needed for a base level description.

Problem Solving Information

This portion of the narrative describes events that interrupt the workflow and management options in dealing with such events. The following shows two examples taken from the M3 narrative.

Part shortages are a recurring situation that interrupts the scheduled workflow. A shortage may occur in the normal build schedule, in which part shortages are identified twelve weeks prior to the kitting date. The appropriate ordering and expediting will occur, but certain kits still may be short of parts when their issue date arrives. A decision must be made whether to hold the kit or to release it as a partial build. The model can be used to identify conditions in which it would be highly likely that parts would "catch up" with a kit that was issued as a partial build. Reliable data would be needed to estimate the arrival time of the parts. That is, ordered parts can be viewed as being located somewhere in the "pipeline" from the vendor to the kit. Milestones of pipeline location can be established and interpreted as delays until part arrival. The decision to reschedule the kit or release it with short parts would be informed by comparing these delays with the expected time for the kit to reach the operation that requires the parts. This is another case in which boundaries are crossed, since estimates of the part arrival date would be provided by Materials Management.

Bumpy loading and alternate board mixes can easily be implemented as parameter changes. However, it is necessary to describe any floor adjustments that are made or could be made in the face of different

STRUCTURAL

- Workstation descriptions: current and proposed.*
- Number of simultaneous jobs that can be performed at a station.*
- Workstation operation: FCFS, Priority, JIT.*
- Board types and routes.*
- Buffer locations: current and proposed.

OPERATIONAL

- Processing times estimates (with set-ups) for various boards and parts at each station. Deterministic and/or probabilistic.*
- Buffer capacities.
- Loading:
  - Issue patterns (kits versus time).*
  - Due dates.
  - Lot sizes (current and proposed).
  - Board mixes for sequential and simultaneous loading.
- Defective boards: classification, percent defective, routing of defectives.
- Learning: initial values of processing times, changes in processing time per board.
- Staffing levels.

PROBLEM SOLVING

- Priority rules when different boards, prototypes, and/or rework vie for a station capacity.
- Blocking rules when buffers are filled.
- Capacity Adjustment options in response to ramping input.
- Rules for allocating cross-trained workers under various loading situations.
- Short Parts:
  - Fraction of kits that are short.
  - Stations at which parts may be damaged
  - Pipeline status of reordered parts.
  - Estimated delays in pipeline.
  - Partial build options with routes and process times
- Down workstations: mean time between failures, repair times, schedule adjustments.

Figure 3 Numerical Parameters and Logical Statements Needed to Specify the Simulation loading and board mixes. This may take the form of rules to allocate cross-trained workers to stations of highest need. Various bumpy load patterns can be explore to identify both appropriate adjustments and the limitations of the cell to accommodate such loads. A bumpy load may occur on a smooth load of A sides, followed by B sides, which would result in a bumpy load from heat sink installation onward. By exploring such options, the model can assist in identifying "good" loading and product mixes.

6. WHAT-IF SITUATIONS AND MODEL DEVELOPMENT

The following will describe how the what-if situations are used to establish model objectives and to guide other model building decisions.

The what-if situations are described within the narratives. In the M1 narrative they appear in general terms, such as, "What are some good options to recover the schedule once a delay has occurred?" Statements like this need a context to be meaningful. As the models evolve, the system is described in further detail. The actual "what if" question may not change, but the context for the question enlarges. The previous discussion of partial builds taken from M3 was absent from M1, even though the what-if question remained the same: "Should a partial build be issued?"

As models evolve, the increasing detail of the narrative serves two functions. The client can develop confidence that the modeler understands the what-if situation, and the modeler can postulate a corresponding technical description of the what-if situation. That is, the situation would eventually be implemented as a parameter change in the coded form of the model. (Here, parameter is taken generally to include structural changes.) If the model structure or level of detail does not include certain parameters, then corresponding what-if situations will be impossible. Thus, as the narratives of the system evolve, they convey the modeler's understanding of the what-if situations to the client, and they force the modeler to postulate levels of detail that can encompass corresponding parameter changes.

Figure 4 shows a summary of what-if situations and corresponding parameters. For each situation, a generic what-if question can be posed as "Investigate the effect upon throughput time, work-in-process and utilization caused by this situation".

This figure is a menu of possible model objectives. It connects "what if" situations with corresponding parameters. Model objectives can be phrased in terms of the ability to address selected what-if situations. Once those situations are identified, the corresponding parameters which must be included in the model are identified. As these parameters will require an appropriate context, they guide the selection of model boundaries, structure, level of detail and required data.

What-If Situation	Parameter Changes
-Introduce Robot for Pick and Place	-Station definition -Process times -Board mix -Buffer location -Fraction of boards needing rework -Partial build options
-Identify appropriate Normal Floor Operation	-Lot sizes -Board mixes -Station capacity -Sequential/simultaneous loads -Buffer locations -Staffing
-Introduce Bumpy Loads	-Issue patterns -Board mixes -Simultaneous/Sequential loads -Rules for allocating of cross-trained workers
-Responses to Part Shortages	-Stations where short parts are identified -Fraction of kits that are short -Pipeline status of reorders -Partial build options -Routes and process times
-Schedule Recovery	-Initial values of WIP -Rules for allocating of cross-trained workers
-Quality Changes	-Inspection process times -Fraction to rework -Routing of defectives -Priority if rework is a previous station
-Capacity Expansion	-Ramping arrival patterns -Strategy for adding capacity -Learning curves for new hires
-Down Stations	-Failure rates -Repair times -Rules for allocating cross-trained workers -Alternate loading

Figure 4: "What If" Situations and Corresponding Parameter Changes

In our study, Figure 4 was used to develop two sets of objectives: One addressed capacity changes in the face of increased loading and the other dealt with managing the floor when the loading and workforce are specified, but the load pattern could be varied and the workforce can be cross-trained to accommodate short term capacity adjustments. This resulted in two models, each with a different level of detail. One is reported in Starr, et. al., 1986

7. CONCLUSION

This work described the process of Structured Feedback as a means of initially developing knowledge of a manufacturing system. The approach engages manufacturing personnel and model builders in a mutual learning experience. It uses the graphical-narrative form of a model as a basis for communication. In the Structured Feedback process, a sequence of models is developed, each being subjected to critical review by manufacturing personnel. The narratives are organized according to figure 1, which provides the "structure". Various benefits of the procedure were pointed out. The role of "what if" situations in the narrative was emphasized; they provide a means for the modeler to gain credibility with the client and they also provide candidate objectives to guide the selection of model structure and level of detail. Use of the approach was described in the context of an electronic assembly workcell. Remarks were organized according to the figure 1 structure and excerpts from a narrative were included. This example, along with figures 1,3 and 4 are offered for others to adapt as needed for similar modeling efforts.

REFERENCES

Shannon, R.E. (1975), System Simulation: The Art and Science, Prentice Hall, Englewood Cliffs, N.J.

Starr, P.J., Skrien, D.S., Meyer, R.M. (1986), "Simulating Schedule Recovery Strategies in Manufacturing Assembly Operations", Proceedings of the 1986 Winter Simulation Conference, J.R. Wilson, et. al., Eds., Washington D.C., Dec. 1986, pp. 694-699.

AUTHOR'S BIOGRAPHIES

PATRICK J. STARR is an Associate Professor in the IEOR Division of the Mechanical Engineering Department at the University of Minnesota. He has developed and taught courses in many topics including Operations Research, System Dynamics, Statistics, Global Modeling, Technology Assessment, Design, and Cellular Manufacturing. He received the Institute of Technology Distinguished Teaching Award in 1973 and was elected Outstanding I.E. Educator by the Minnesota Student Chapter of I.I.E. in 1982. His consulting and research has included model building in a variety of settings, from the definition and analysis of solid waste management options to the flow of circuit boards through various processes. He has published in a variety of journals, and is currently doing work in sensitivity analysis of large scale systems, manufacturing simulation and expert system.

THOMAS M. KEENAN is an Applications Engineer for the Governmant Systems Group of Control Data Corporation located in Bloomington, Minnesota. His duties include the application of statistical and operations research methods to solve manufacturing problems. He received a Bachelor of Statistics degree from the University of Minnesota in 1986. He is a member of the Institute of Management Sciences.