**New Directions in Modelling
Production and Distribution Problems**

by

Uday S. Karmarkar

CMOM 94-01

**CENTER FOR
MANUFACTURING &
OPERATIONS MANAGEMENT**

Working Paper Series

1994

WILLIAM
SIMON
GRADUATE SCHOOL
OF BUSINESS ADMINISTRATION
UNIVERSITY OF ROCHESTER
ROCHESTER, NEW YORK

New Directions in Modelling Production and Distribution Problems

Uday S. Karmarkar, William E. Simon Graduate School of Business Administration , University of Rochester, Rochester NY 14627, USA

Summary: Production and distribution problems have historically been a major category of research in Operations Research. Yet, applications of OR/MS methods in production and distribution practice remain limited and OR/MS groups are not thriving in industry. The reason may be in the way problems are approached in practice, which then has consequences for modelling and solution methods. Perhaps a slightly different perspective for OR/MS could lead to significantly greater successes in applications to production and distribution problems.

Introduction

Work on production and distribution problems remains the largest category of publication and research in Operations Research and Management Science (OR/MS). A very substantial proportion of all the papers submitted and published in OR/MS related journals deal with these topics. The application of OR/MS like methods to production problems could perhaps be said to start with the Economic Order Quantity model of Harris, Wilson and Camp more than seventy years ago, when the term OR/MS did not yet exist. The subsequent years of research on production and distribution problems have produced some successful applications, and a great number of research papers. Unfortunately, too many of the latter cannot be described as having had any impact on the practice of manufacturing. Apart from the low effectiveness of formal work, it is also apparent that practitioners in manufacturing and distribution environments do not necessarily turn first to OR/MS analysts to find answers to their problems. The most widely applied formal techniques in production control are Material Requirement Planning (MRP) and simple release rules, that do not make very much use of OR/MS techniques, modelling or optimization. Distribution management has been a somewhat better arena for OR/MS models. Many commercial distribution management software packages and practical applications do make use of modelling techniques. However, even here, the proportion of research output that finds actual application is small.

I do not mean to suggest that OR/MS has little to offer to practice. On the contrary, the opportunities for OR/MS applications are greater than ever due to the twin trends of the increasing costs of management and the ubiquitousness of cheap computing power. Nor am I arguing for a cessation of formal theoretical

research. Clearly, such research is a part of the "supply chain" in that it provides the groundwork for applied techniques. However, I believe that the potential for application of OR/MS methods is not being fulfilled. The reality is that in industry as well as in academia, the role for OR/MS is declining. Many firms have closed down their OR/MS groups completely. There are fewer departments in academia, and courses on these topics are less popular, noticeably in schools of management.

What then of the opportunities for modelling and analysis due to cheap computing and expensive labor. The issue is that these opportunities will be garnered by other disciplines such as Artificial Intelligence (AI) and Knowledge-based Systems (KBS), or by spreadsheet and simulation models. While these methods ultimately are not really that different from OR/MS, it appears that many valuable opportunities are being lost to the field. Part of the reason for this is that OR/MS education and research has tended to emphasize formal academic research and rigorous publication standards at the expense of good applications. Many well-known authorities have already inveighed against excessive formalism and theoretical research in OR/MS and I will not add more noise to that clamor. Rather, I propose to provide some observations regarding the obstacles encountered in applied work, and some practical suggestions on making progress in that area. Let me also hasten to admit that my qualifications to pronounce on these issues come as much from failures as from successes.

In the following, I first identify some broad inconsistencies between standard OR/MS paradigms and the needs of practice. I then discuss two cases from my own experience that illustrate some of issues I have raised. Next, I address in turn the issues of modelling approaches, solution strategies and solution methods. Finally, I list what I perceive as some of the outstanding opportunities for OR/MS and modelling methods in production and distribution management. These necessarily reflect my own perceptions, and there are undoubtedly many others that have not occurred to me.

The Limits of OR/MS Paradigms

Perhaps the most prevalent paradigm of OR/MS is that of the optimization model. In the paradigm, decision problems are phrased in terms of the familiar form of the extremization of an objective function, subject to a set of constraints on the decision variables. Constraint sets are the main tool for describing the context of the decision problem. One of the major contributions of OR/MS has been the identification of certain standard optimization problems, together with the construction of algorithms that can solve these

problems to optimality. Outstanding examples include LP and the Dantzig and Karmarkar algorithms; Integer Linear Programming and Gomory's cutting plane algorithm and branch-and-bound; and various network flow and graph theoretic problems. One consequence of these successes has been a strong bent in the profession towards thinking in terms of abstract optimization problems and provably optimal algorithms. Yet the needs of practice can be quite different. Even the basic notion of optimality can really only make sense in the context of a formal model. Given that any such formal model is typically a limited and approximate representation of reality, to talk of optimal solutions to a real problem is often not meaningful.

The optimization formalism tends to foster a world view in which abstract models become a framework for representing real problems. The next step is to develop an algorithm which has desirable properties of convergence and optimality. The task of application and problem solution is then seen as a rather mechanical exercise of acquiring data and processing it through the available machinery. Implementation of the solution is even less interesting from this perspective. Practitioners however, do not view the world through the optimization framework. What they tend to see are streams of events and observations that are recorded as data. Practitioners also act within the context of management processes that are triggered by these events and observations. The resulting data stream is the primary view of the world for practitioners. Schematically, the OR/MS paradigm might be represented as:

World ----> Model ---> Algorithm ---> (+data) ---> Solution --> (implement?)

On the other hand the natural perspective for the practitioner might be more like:

World ----> Data ---> Analysis ----> Process ----> (model?) ---> Resolution

Perhaps the differences have been emphasized here more than the similarities, but the role of data and implementation in the former paradigm is underplayed as much as the role of models and algorithms is in the latter.

It might help to see the differences more clearly if these schemes are examined as models for a first course in OR/MS. The typical OR/MS course in a management school tends to follow the first scheme. The standard formulations are quickly introduced as templates for real problems. This is followed by algorithms with toy examples to illustrate the models and algorithms. Real world data is rarely seen. But imagine a

course designed around the second scheme. It would present the students at the start of the course with data sets (numerical and otherwise) drawn from real situations. The students would first be taught to organize and analyze the data. Then they would be taught how to extract descriptive models of relationships, dynamics and feasibility. The course might describe how data and information drive management processes. Now models could be introduced as powerful tools to structure the decision points in the management process. The approximate nature of models would be very clear at this point, and concerns regarding algorithmic convergence and optimality would perhaps be muted.

The Primacy of Data and Process

I will hereafter use the terms "information" and "data" interchangeably simply to avoid a discussion of the differences. Information availability is a key limitation in developing solutions. It is true that specifying information requirements is part of the modelling process. However, it is a useful discipline to start the modelling process by first modelling existing information flows and processes. If these are so ineffective as to require re-engineering of business processes, then modelling is premature. Reengineering and simplification of business processes is generally a pre-requisite for modelling and building decision systems.

It should be made clear that when I refer to modelling information and data flows, I do not mean the development of data structures for algorithms, or even the organization of data base management systems (DBMS). While the latter may be more relevant, the real issue is to understand the interplay between information flows and management processes. Some critical information may not be (as yet) captured in formal DBM systems, and simply identifying data and information flows will not convey a sense of the management process by which actions and decisions are actually executed. Thus, a shop data collection system, and the associated data bases cannot convey the process of order management on the shop floor. The structure of an MRP system does not determine the management process by which the system is actually used to develop a consensus plan between marketing and manufacturing functions. The latter agreement requires a sequence of meetings and "sign-offs" which is usually very clearly defined in terms of process steps as well as responsibilities. A failure to understand these management processes will doom any modelling effort aimed at application of the results.

The practitioner's emphasis on data, information and process leads to significant differences with the

world-view of the OR/MS specialist. My perception of some of the differences is described in Table I.

The view in practice reflects the realities of management and decision processes, where communication, common understanding, agreement or consensus, responsibility and clear action plans are crucial issues. The OR/MS world on the other hand, tends to be driven by structure and rigor.

The treatment of time is a good example of some of the distinctions in Table I. OR/MS often treat capacity in terms of loading time buckets. However, in practice, it is much easier to think in terms of time

Table I: A comparison of different modelling perspectives in OR/MS and practice

<u>OR/MS</u>	<u>Practice</u>
Optimization	Feasible improvement
Single best solution	Many good solutions
Single agent/objective	Multi agent/objective
Convergence = optimality	Convergence = consensus
Decision making	Management process
Standard engines (e.g., LP)	Robust computation
Structure oriented	Data oriented
Homogeneous structure	Mix of different structures
Abstract variables	Actual events, objects
Decision oriented	Action/event oriented
Formal model	Virtual representation

lines and events and intervals, much as in a calendar, diary or simulation model. For example, if we were decomposing a process across several stages in practice, we would identify the decomposition in terms of points in time at which certain events took place. Thus a trip might be thought of as arrivals and departures at locations, with the intervals to be scheduled. Similarly, delivery schedules and lead time intervals are characteristic of production and distribution co-ordination. In some models, such as scheduling with Gantt charts, we use this kind of modelling, but we lack ways of dealing with decomposition of composition of these models. As a result, capacity and planning models are often formulated very differently. Perhaps this

is one reason that time interval and release oriented methods like MRP and DRP are used in practice for planning even though they are completely unable to actually deal with resource allocation decisions.

To reiterate, focussing on information flows and management processes, can create an entirely different view of modelling and solution methods. The emphasis shifts from rigor and homogeneous structure, to accurate representation and robust feasible solutions which are acceptable to all parties. The case examples of the following sections may serve to illustrate.

A Production Scheduling Example

The first of my examples is drawn from the development of a large scale scheduling system (Dobson et al., 1992). This system represented a deliberate attempt to span the range of OR/MS techniques from theoretical development of models and algorithms to a practical implementation of models and heuristics that could be used to solve large realistic problems drawn from actual factory data sets. In this effort we discovered that while we indeed did many things right, we did not go far enough, at least in our initial design.

Among the things we did right were:

- The development of new theory to address real issues. The technical scheduling literature, while vast, offers little useful guidance on issues such as scheduling with multiple objectives. We were able to make new contributions in the area of understanding the novel issues of combining batching and sequencing, in scheduling parallel heterogeneous machines, and in scheduling resources that are required simultaneously (Karmarkar, 1987; Dobson et al., 1987, 1989). Work continues on other such models such as scheduling, batching and loading with batch processors (such as furnaces or kilns), and scheduling with finite buffer capacity.
- The development of "plug compatible" models. One of the problems with much of scheduling research is that very few results are amenable to combination with others. We emphasized the development of heuristics that worked in a common compatible framework, even at the expense of giving up on rigor and optimality. Our approach is based on concepts from Lagrangean decomposition (Dobson and Karmarkar, 1986).
- The scheduling system was designed to interface with MRP systems and with shop floor data collection systems. The data models used conformed to standard practice, but were extensible to enable richer

modelling of specific processes.

- The system can work with incomplete data on costs and objectives, and objectives can be flexibly specified to include costs, and priorities.
- The system has some high level "knobs" or parameters which can be used to change the behavior of heuristics at an aggregate level, and which can be used to generate alternative solutions and scenarios. Internally, these settings are used to change heuristic parameters and relative weights on conflicting objectives such as capacity utilization and flow times or work-in-process.
- Optimality is sacrificed for speed. The system is extremely fast but is nevertheless able to significantly outperform other practical systems in tests with actual factory data (Dobson et al. 1992).
- As a "proof-of-concept", the system has been interfaced to a real-time shop floor system that also connects the system to certain MRP systems.

Despite these efforts to be realistic we have discovered that we need to go much further:

- We assumed that a standard DBMS would suffice for data management. It turns out that this approach is too slow, and we had to develop special procedures to work in a real time environment.
- Certain data elements that we assumed would be available in most plants, were not there in a plant that justifiably claimed to be the most data rich in the US.
- There were some standard fabrication processes that we could not model compatibly within our range of scheduling models.
- Certain industry practices in resource loading and scheduling could not be modelled even within our heuristic decision framework, and we doubt our ability to do this in the future.
- Practitioners felt much more comfortable with a simulation model, and in fact could not really distinguish between a feasible simulation run and an optimization approach. Fortunately, our scheduling model could essentially simulate as well.
- The scheduler solved difficult problems that practitioners did not think were feasible. However, on simple problems the solutions often did not look simple, and were hard to explain. Trying for every possible improvement sometimes led to complicated schedules that the practitioners did not understand.
- Practitioners were often more concerned with the user interface and reporting tools than the quality of the solutions produced.
- The system did not have the ability to refine and improve a user defined schedule.

The next version of this scheduling system will be an object oriented, simulation like model that will

be able to utilize multiple solution techniques from OR/MS and AI/KBS. The system will accommodate "best practice" constraints and user defined rules. Any schedule, whether user or system generated, will be treated as a proposal that could be refined or modified through improvement heuristics and constraint satisfaction techniques. Graphical interfaces will be integral to the system. Nevertheless, I am sure that we will discover a number of new limitations to our approach.

A Distribution System Example

This "example" is really a composite of several different experiences with seemingly different distribution environments. They included a national distributor of industrial equipment, a manufacturer and national distributor of household cleaners, and a manufacturer and global distributor of medical products. In each case, the issue was the redesign of distribution systems including location and role of facilities, and transportation policies.

We discovered rapidly that we were unable to use traditional modelling methods for several reasons:

- These firms were facing competitive markets and were not interested in simply cost minimization.
- The distribution system design thus required significant mediation between marketing, distribution and manufacturing functions addressing basic company strategy issues.
- Changing the system entailed different contractual and business relationships with other entities in the distribution channels including independent distributors, freight forwarders, and agents.
- Strategic positioning was often defined in terms of a service-time profile that was affected not only by location, but also by inventory deployment and transportation policy. All these therefore had to be addressed.
- The roles of locations changed and had to be clearly defined. Some of these roles were: consolidation, breakbulk, seasonal storage, JIT warehouse, backup shipment warehouse, etc.

The net result was that we used simple heuristics, spreadsheet analyses and simulation to construct and evaluate a small set of design options within the limited time available to us. This was a stop-gap solution which nevertheless did the job, especially in terms of strategic positioning and revenue vs. cost tradeoffs. Our longer term aim is to develop models and solution methods that can address these complex problems in a way that communicates with practitioners.

Our approach at present is to use traditional methods to formulate the problem, and to analyze very

small instances (Saxena, 1993). The formulation includes costs as well as service time requirements. The decision variables include locations, shipment paths, ship modes and frequencies, and inventory deployment including seasonal movements and safety stocks. To solve realistic versions of the problem, we are developing a combination of genetic algorithms and heuristics (Karmarkar and Saxena, 1992).

So far we have found that on very small problems, with 3 items, 3-6 retailers, and 2-6 potential sites, the genetic algorithms (GA's) produce optimal solutions in over 85% of the examples tried, with small gaps in other cases. In somewhat larger problems, with 4 supplying plants, 2-10 items, 8-12 retailers, and 12 potential warehouse and DC sites, the GA solutions average within 8-10% of the best lower bound. These are non-linear problems with up to 2500 constraints, 280 0-1 variables and 2000 continuous variables. We cannot presently compute optimal solutions.

We feel confident that GA's are a reasonable tool for large realistic problems, but we will have not ability (bounds or otherwise) to benchmark the quality of the solutions. We also do not have any other alternatives. We still have some way to go on modelling competitive situations, and on incorporating issues related to business relationships, contracting and facility roles.

Modelling and Representation

Examples like the above – and I am sure the reader has encountered similar situations -- lead me to reconsider our modelling approaches. There has been a tendency in OR/MS to identify homogeneous standardized modelling structure, and then to force problems into them. The problem is that real applications do not categorize that easily. The structured approach tends to be limiting in terms of the objectives as much as the constraints. There is an observable tendency to cast problems in terms of cost minimization. There are certainly many problems that can be cast purely in terms of cost. However, there are many that cannot. The current interest in industry for programs such as JIT and TQM might be interpreted as a recognition that the traditional cost and revenue decomposition of objectives is proving to be inadequate for today's competitive markets. Many companies in the US have made real shifts in their management emphasis from traditional purely financial measures to issues of lead time, quality and customer service. It does not appear that we have kept up in terms of model development. It is worth noting that after many years of research, few of our models address the time and quality impacts of decisions. It is only in the last five or ten years that we see a few models of production and inventory management decisions capable of analyzing time

related issues, and I have reviewed some of this work in Karmarkar (1992). The situation with respect to other objectives such as quality and flexibility is even poorer, though I cannot claim to have authoritative knowledge of its current state.

An analogous situation exists with respect to the modelling of constraints and process dynamics. Since I do not have any ready made answers, let me at least pose some questions:

- *Do we have any way of modelling existing data and information flows?* There are information systems methodologies such as Data Flow Diagrams (DFD) to record and represent such flows but they do not have a symbolic representation, and in fact have no special methods for capturing decision points. Control theory methods have tended to stay with real time execution and control, and have not reached application to resource allocation problems. Petri nets may still have some promise in this regard, but do not seem to have accumulated analysis and solution techniques to accompany the representational method.
- *Can we model existing procedures and business processes?* MRP is one obvious example that I have mentioned earlier. There are indeed some models of MRP that can be included in decision formulations. There are increasingly effective models of production release procedures such as Kanban. However, we do not have an understanding of many key management processes even at the level of a taxonomy. For example, how are marketing and manufacturing needs reconciled in issues such as making forecasts, setting production and sales targets, or setting quality levels? How are process and communication issues to be included in multi-stage multi-facility manufacturing and distribution management?
- *Can we develop models of multi-agent decisions?* Game theory gives us some tools to represent both co-operative multi-agent decisions as well as non-cooperative or competitive situations. However, the methods though rigorous, are limited to simple examples. We need to be able to tackle issues such as contracting in manufacturing networks, in sourcing and supply chains and in distribution channels. Such analyses today are too often completely disjoint from the underlying production and inventory models, which in turn usually do not consider these aspects.
- *How do we handle issues of competition?* Consider that the location literature is dominated by cost minimization formulations. Yet in many service industry examples, location is a competitive exercise, since customers choose the sites they will visit for service. Similarly, product design, service response, quality levels are all subject to competitive action. We often formulate models that assume that recourse can be made to subcontracting for say, transportation vehicles. However, it is often the case that such

demands are seasonal and result in competition among demanders for scarce resources. Similarly, acquiring capacity or technology, as with a large flexible manufacturing system, may have to be analyzed in the light of competitors' responses. Again in addition to game theory, there are economic models that address these issues, but largely at a conceptual level.

- *Can we build modular, virtual, plug compatible models?* In a sense, an LP provides a framework for modular modelling, since many different aspects of a firm can be represented in an LP model. However, the linear structure that provides the standardization is too limiting for many problems. Perhaps formal methods of decomposition can suggest such modelling approaches. It becomes apparent that part of the problem is in the limitations of solution methods that require a very homogeneous structure to work. As examples of modular schemes let me point to queuing networks of both the open and closed variety, and Petri nets. We do not have an analogous methodology in say, scheduling or planning models. The ideal would be close to simulation in its modularity: it is relatively easy to add more to a simulation or to change the model without affecting other aspects; furthermore, the analysis process is unaffected by such changes. The process of representation is transparent to the practitioner and as a result the model has high credibility. Of course, simulation is a lot easier than decision making, but perhaps there is some common ground. Again, the availability of solution techniques is likely to be the limiting factors
- *Can we accommodate "best practice" requirements, and heuristics?* This is relatively easy to do in terms of constraints, except that the structure of the constraints may not be very convenient. So-called expert systems are a natural way of modelling such situations, where a method, guideline or practice has been found to work by trial and error over years or even decades and centuries. It is typically difficult or impossible to extract and state the underlying optimization problem. Such examples are common in process industries.
- *Can we model learning, and process improvement?* In the past, a typical shop floor control system would have focussed largely on material tracking and job scheduling. Today, quality control and process improvement are likely to be as if not more important. We have many models that address the evolution of systems in terms of production scheduling events and material movement. We are much poorer in terms of models of process improvement and learning. Our models tend to treat these phenomena as static, with given probability distributions. In fact, knowledge about processes changes over time, and the rate depends on how processes are managed. That is to say, process scheduling and process improvement interact with each other. Just for example: shorter production lead times aid learning and

process improvement. On the other hand, complex scheduling of processes with many changes, may obscure cause and effect relationships, and hinder learning. In such areas we are weak even in terms of simulation methods.

I do not think that the questions above are that difficult to address. However, for whatever reason, OR/MS has not actively pursued such modelling approaches. I would argue that as one becomes interested in applications, and starts with the observation of real world problems, processes and data, the need for such methods will be apparent very quickly. On the other hand, if the direction in OR/MS is towards abstract research, such developments will not seem attractive because of the difficulty of solving problems optimally or establishing analytical results definitively. In other words, the difficulties of rigorous analysis and solution may have been inhibiting factors in the past.

Solution Strategies and Methods

By now my drift with respect to solution techniques should be apparent. Some of the great successes of OR/MS have come in the area of optimal solution techniques for significant problem classes. Progress in many areas continues to occur, and indeed there are many problems which can be modelled and solved by these methods. However, we have also come to recognize that there are also large classes of problem where optimal solutions are very difficult to obtain, even assuming large improvements in computing capability. Many of the best opportunities open for OR/MS fall in this latter category; most of the standard homogeneous problems have probably been identified already. It is likely that we will hereafter increasingly be in an environment not unlike that of the PC: no sooner is an application available than there is a demand for an extension or an enhancement. The most successful applications will be robust, easy to use, easy to modify and easy to build upon.

If our models must be flexible, extensible, and modular, we need appropriate solution strategies and solution methods. If we are willing to give up on rigor and optimality to some extent, many useful techniques already exist. Let me reiterate here that rigor remains important in basic research. However, in real problems where models are often poor approximations to reality, rigor and optimality can be misplaced objectives.

There have been many solution strategies and techniques that have appeared recently from the domain of artificial intelligence and knowledge based systems. Many of these techniques are not dissimilar to

traditional OR methods. However, they tend to be less concerned with rigor, often because they were created to deal with extremely difficult problems such as chess playing, theorem proving and pattern recognition. They include search heuristics such as simulated annealing, tabu search and beam search, which are very closely allied to traditional gradient search methods and tree search methods. Other techniques like genetic algorithms, neural nets, blackboard methods, constraint satisfaction, learning systems, and production systems or expert systems may seem further afield. In these cases too, I would argue that the differences and novelties are much less than the similarities and in fact there are significant gains to be had from combining techniques. This is not the appropriate place to delve into these opportunities however, and I will limit myself to addressing what we might expect in the specific area of production and distribution modelling.

Many of these techniques are already being widely used, as a glance at the research papers in this conference shows. However, there is still a tendency to think of a model as having one type of structure with one solution technique. Alternative approaches are seen as competing rather than complementary. I would assert that this is not the right strategy to accompany a robust modular modelling approach. Rather I would suggest the following:

- *There may be many alternative acceptable solutions, rather than one optimum solution.* And it may require different methods to generate them. Some methods such as genetic algorithms naturally generate and maintain multiple solutions.
- *Any "solution" should be treated as a proposal to be changed or improved.* Even realistic models fall short of reality. Modifications to solutions will be necessary as new constraints and objectives are identified. It should be an intrinsic part of a solution strategy that solutions can be constructed and/or improved. The solutions may come from algorithms or from a user.
- *An implication is that constructive methods must be supported by improvement methods and heuristics.* Traditional methods are especially weak on improvement. Even current improvement methods such as exchange heuristics can be too local and limited. For example, it may be required that a solution to a scheduling problem be modified to be more intuitively acceptable, or to cleave to some practical guideline.
- *Expert systems may provide methods for accommodating best practice modelling.* Expert systems may also provide methods for improving starting solutions obtained by more structured methods, to then meet complex constraints on the form of the solution.

- *Alternative solution techniques can gain by trading solutions.* For example, simulated annealing can be used to enhance solutions generated by genetic algorithms. The results from GA's can be used as starting points for other methods.
- *There must be a mechanism for transferring information between different methods.* In my opinion, the GA framework can serve as a central "blackboard" or clearing house for alternative solutions generated by different methods.
- *Parallel computation will be an essential part of solution strategy.* Future computing advances are likely to rest on parallel computing. We need to develop intrinsically parallel methods in addition to those already available such as production systems and genetic algorithms.
- *General algorithmic engines could be equipped to learn about specific problem domains.* Techniques such as neural nets and genetic algorithms have been used to improve solution heuristics by learning about specific parameter domains that occur in particular applications.
- *We need a common language for models that can be understood by alternative solution techniques.* Current efforts in developing modelling languages are still aimed at a math programming framework. It should actually be simpler to develop a language that will be understandable by robust techniques.

Opportunities for the Future

I reiterate that the trends of increasing labor costs and cheap computing present the OR/MS professional with an unparalleled opportunity. Companies are trying as never before to improve and automate their management processes and to increase the productivity of personnel in all areas. However, the opportunity could be lost, if OR/MS remains limited to traditional decision modelling. This was perhaps appropriate 20 years ago, but is not in today's environment. It is necessary to recognize that to have an impact on applications, we need a paradigm that is founded in modelling at the level of data acquisition and management process. Furthermore, we also need modelling methods that are as robust and modular as simulations, spreadsheets or other modern PC applications. Changes in both modelling and solution strategies appear to be necessary.

Let me briefly list some of the application opportunities and trends that I perceive; I am sure the reader will be able to list many more:

- Applications are moving from generic models to specific areas. So-called *vertical applications* are the

prevailing trend. That is to say, production control systems or vehicular routing models are focussing on particular industries rather than on generic methods. A scheduling system must address say metal cutting or chemical batch processes, instead of claiming to be a general purpose tool. It is already very clear that the models that can realistically solve one kind of problem can be worthless on another, because of the specialized heuristics, model structures and management processes that are used. Interestingly, commercial software packages also suffer from the problem of being too generic. MRP and scheduling systems are a good example where many of these systems are still sold as though they can work for any plant. A few software companies are moving rapidly to exploit this opportunity by creating differentiated products.

- From the research viewpoint, there are some very interesting problems in modelling specific manufacturing and distribution environments. Some examples are:

Hot metal processing (e.g. steel casting and hot rolling)

Batch chemical processing

Semi-continuous chemical processes.

In each of these areas, there are differences across plants and products.

- The entire subject of managing operations for *profitability* in a *competitive* environment has been little addressed in either the OR/MS or any other literature. Just for example, capacity decisions or location decisions are significantly different in a competitive framework than in a cost minimization, single firm demand independent framework. In our traditional models, the shadow price of capacity only depends on internal factors. In the real world, if a competitor builds capacity, the value of every other company's capacity drops. I have already mentioned other examples such as location models and the need to look at objectives such as service, time response, innovation etc.
- There is a need to develop solution engines that can utilize a variety of algorithms from OR/MS and AI/KBS. The first issue to be addressed is probably the development of a suitable architecture in which multiple techniques can be accommodated in parallel.
- There is a great deal to be gained from combining ideas from different areas. For example, current GA techniques are weak on maintaining feasibility with complex constraints. Developing GA's that can exploit problem structure would be extremely useful in improving performance, while maintaining robustness.

It may seem that it is easier to talk about the issues that I have raised in this article, than to actually do

something. However, I would argue that many research groups are pursuing every topic that I have mentioned. To emphasize my own interests in these topics let me end with a brief list what I am doing along these lines in addition to the two examples given earlier. I am working (with R. Jaikumar, Pierre Lefrancois, Benoit Montreuil) on the development of a shop floor system that will integrate data collection, simulation and decision support, for both material management and process learning or improvement. The aim is to design a platform that will support a variety of decision support tools, which might be developed by others in the future. In the area of scheduling, I would like to explore the differences between batch chemical processes and metal fabrication, and to develop modular modelling techniques that work like simulation modelling. I have already mentioned some ongoing work in distribution systems. I invite the reader to correspond with me either to share ideas or experience in any and all of these areas.

References

- Dobson, G. and U. S. Karmarkar (1986). "Large Scale Shop Scheduling: Formulation and Decomposition". W. E. Simon Graduate School of Business Administration, Working Paper Series No. QM 86-31, University of Rochester. To appear in *Optimization Models and Concepts in Production Management*, Villa A. and P. Brandimarte (eds.).
- Dobson, G., U. S. Karmarkar and J. L. Rummel (1987). "Batching to Minimize Flow Time on One Machine". *Management Science*, **33**, 784-799.
- Dobson, G., U. S. Karmarkar and J. L. Rummel (1989). "Batching to Minimize Flow Times on Parallel Heterogeneous Machines". *Management Science*, **35**:5, 607-613.
- Dobson, G., U. S. Karmarkar and J. L. Rummel (1992). "A Closed Loop Automatic Scheduling System (CLASS)". *Production Planning and Control*, **3**:2, 130-140.
- Karmarkar, U. S. (1987). "Lot-Sizing and Sequencing Delays". *Management Science*, **33**:3, 419-423.
- Karmarkar, U. S. and R. Saxena (1993). "Distribution Systems Design Using Genetic Algorithms". William E. Simon Graduate School of Business Administration, Working Paper CMOM 91-03, University of Rochester.
- Saxena, R. (1993). "Multi-Item, Multi-Echelon Distribution Systems Design". Ph.D. Thesis, William E. Simon Graduate School of Business Administration, University of Rochester.