Mathematical Programming 42 (1988) 99-112
North-Holland

APPLICATIONS OF OR IN THE OIL INDUSTRY

B.R.R. BUTLER
The British Petroleum Company p.l.c., Britannic House, Moor Lane, London EC2Y 9BU, UK

Received December 1987

Key words: Modelling, oil production, operational research, pipelines, scheduling.

Introduction

When I chose the title for this talk, I hadn't quite appreciated what a lot of ground it implied, so I think that it would be better if I limited myself principally to discussion of the application of Operational Research to the Exploration and Production phases of the oil business (what we refer to as the "upstream side").

I will however briefly mention some of the applications of Operational Research in the "downstream" side of the business (supply, refining, marketing).

Before I start to discuss applications, I should establish the background for my talk by briefly describing the nature of the "upstream" side of the oil business and mention my views on the nature of Operational Research.

In order to be able to make profits from the production of oil and gas we first have to identify areas where the geology is suitable and where the political conditions will permit exploration, and where geographical conditions would permit the successful development and operation of any discoveries we make.

Having gained access to such regions we must explore using geological and geophysical methods to identify sites where hydrocarbons may have accumulated, and decide whether the probability of success and the potential size of the accumulations justify drilling.

If the drilling of exploration wells discovers oil or gas, we carry out further delineation drilling and geophysical prospecting to accumulate enough data for us to decide on a development programme.

If the discovery is offshore or in a remote area the decision to go ahead with development commits us to very large expenditures (c. £1 billion in the North Sea) and often to the solution of difficult technical problems.

Having developed a field we must monitor its performance and exercise control so as to make the most of what we have found.

Always, of course, we have to be able to sell what we produce profitably, whether this be crude oil or gas sold directly to others or refined products marketed by BP.

Always we must be aware of the effects of our activities on the environment and the importance of people—our own employees and our hosts and neighbours.

In addition to the other problems, our decisions must be made in the light of the many uncertainties involved and, in particular, since we are faced with very large "up front" expenditures, a great deal of economic uncertainty—our rewards come over a long time span, and we cannot be at all certain about oil prices, or exchange rates, or competitive activities in the future.

I hope that these few words will give some inkling of the range and complexity of our operations.

Now, what do I mean by Operational Research?

I believe that a reasonable definition is that it is the application of scientific methods to solve business and technical problems.

It may involve many disciplines or specializations, and I think it should be regarded more as an attitude of mind than as a specific set of techniques.

Although I emphasize the "scientific" approach, I think that there is an element of "art" involved in judging how much of the real world and how much detail should be considered in order to obtain a valid solution to a particular problem.

When speaking of the scientific method and the element of art involved perhaps we can recall some remarks of Martin Beale about the "Scientific Method" [1]. This is traditionally viewed as consisting of the stages:

1. Observation,
2. Hypothesis,
3. Prediction,
4. Experimentation,
5. Verification.

Beale, quoting Sir Karl Popper [2], observed that the way science and OR is really done consists of three stages:

1. Problem identification,
2. Trial solution,
3. Error elimination.

This view suggests that the model comes first, before the data, and repetition of stages 2 and 3 helps to define an appropriate model and the data relevant to the problem.

The development of the model and collection of the data form a very important means of communication between the various interested parties; management, operations, technical specialists and model builders, and it is important that they are all involved. The purpose of the model is to integrate their knowledge, skill and efforts.

When a model is developed for regular or routine use it is very necessary to ensure that the "error elimination" stage is kept up and that the communication and integration of effort continue so that the model remains relevant and that everyone remains involved and committed to it.

Usually there is no unique solution to a problem, and then we seek a "best" solution. We have of course to know by what criteria we judge best. We would like to be able always to say "most profitable" but this is not so easy as it sounds, since our information is always uncertain. Particularly important areas of uncertainty are the predicted political and economic environments. There is also the difficulty of definining "marginal costs and benefits" with any accuracy.

Applications

I will now try to describe some applications in which I have been involved [3] or have come across in my career. I will try to be brief and outline the main features of each problem, the solution techniques used, and the success or otherwise achieved.

The first problem was concerned with the day-to-day control of the production of the oil-fields in Kuwait.

There had always been the need to select which wells to produce to meet the daily production target, keeping the quality of the oil produced within constraints. A simplified crude oil production network is shown in Figure 1.

By the mid 1960's this problem became complicated because of:

(a) Increasing utilization of the gas evolved from solution in the oil, including the use of gas for re-injection to maintain reservoir pressure.
(b) The development of a bulk Liquified Petroleum Gas ("LPG") operation.
(c) The development of new oilfields.
(d) The increasing awareness of the need to control the distribution of production within reservoirs to avoid premature local water breakthrough leading to increased costs and reduced total recovery.

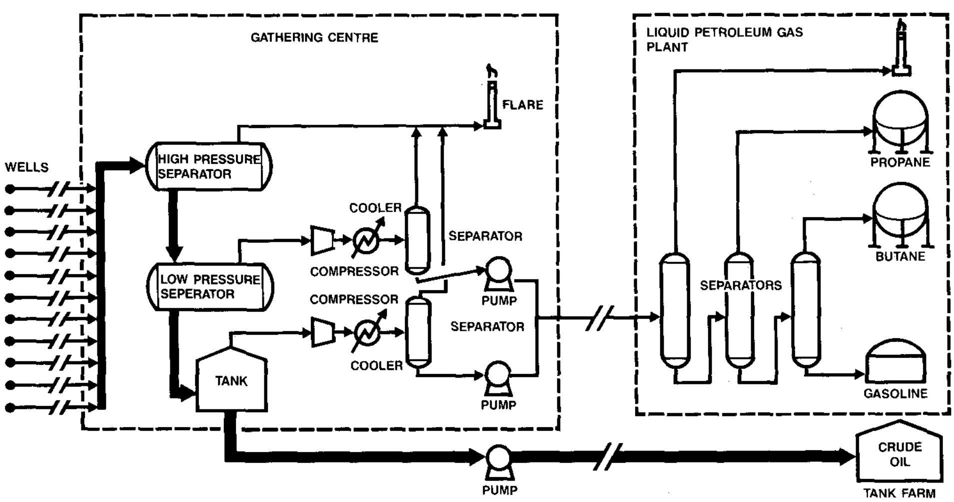

Fig. 1. Simplified crude oil production network.

At the suggestion of Dr. J.E. Warren, a linear programming model was developed in 1964 to solve the problem. This became known as the "Selective Production Scheduling" model (SPS) [3, 4].

It was not very easy to value any solution directly in money terms, so we used a composite objective function:

—to minimize the shortfall in meeting oil and gas demands,

—to maximize the production of LPG,

—to minimize the unusuable volume of gas produced, which had to be flared,

—to match as closely as possible the long term pattern of production within and between reservoirs to ensure long term benefits.

The principal constraints were on:

—crude oil quality,

—plant and pipeline capacities,

—well capacities.

Originally solutions to the problem were run in London, to cover a range of production rates.

Amending these solutions to accommodate to actual well and plant conditions was not easy. So when suitable computers became available on site the problem was run in Kuwait on a day-to-day basis.

It was necessary to develop a Production Information and Control system to hold up to date well and plant data. A Matrix Generator programme, which generated a matrix involving 300–500 rows and 1300–1500 variables, was fed from this; the LP was solved and reports generated; the eventual version sent lists of wells to be opened and closed direct to the operations in the field.

It's hard to prove that this system actually saved money, but it did provide a reliable and consistent method of resolving conflicting requirements. Originally treated with some scepticism by operations personnel in the field, in latter days there was some consternation if, as occasionally happened, the computer system went down.

A problem that this system could not solve was what the daily overall production target should be.

Most of the oil produced was exported in tankers. There was a good deal of information available about the forthcoming oil lifting pattern, and information from the tankers themselves as they approached, but because of delays, diversions, changes in programme and weather there was some uncertainty about actual offtakes. A large tank farm (storage facility) acted as a buffer but it could not absorb all the variation; the production rate had therefore to be controlled. Allowing the storage to fill would mean a complete shutdown, leading to a lack of gas produced, along with the oil, which was used for power generation and water distillation. Allowing it to empty would mean that tankers were kept waiting and a need to go onto maximum production disrupting scheduled maintenance programmes.

When this problem was put to Martin Beale, he immediately suggested a Dynamic Programming solution [3].

A mere five hundred pounds worth of programming later, the essential algorithm was programmed and working.

Again we used a composite objective function being the weighted sum of the three sorts of penalties shown in Figure 2:

(1) A penalty based on the predicted stock level at the end of each of the next seven days,

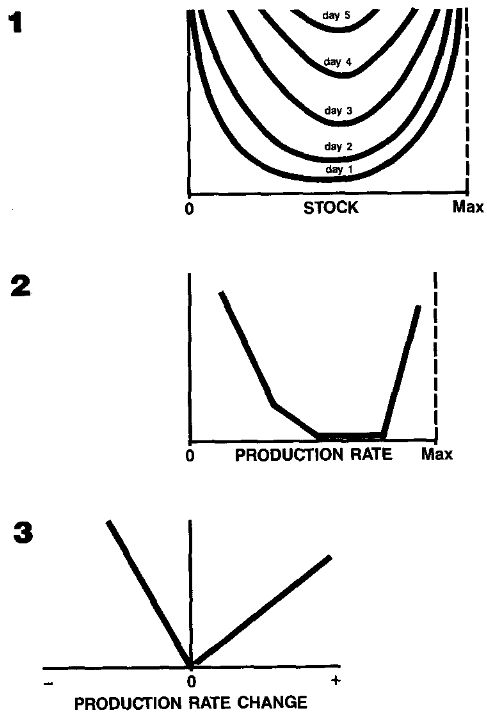

Fig. 2. Penalty functions used in production rate settings.

(2) A penalty based on each daily production rate,

(3) A penalty for changes in rate.

From historical data we were able to fit predictive equations relating actual offtakes to shipping information and find the variance of the prediction from the actual offtakes.

For any assumed production schedule we would predict the stock level at a succession of time steps. The error in the prediction of the stock levels would be the same as the error in the prediction of offtakes. We based our estimates of these errors on historical experience. We could therefore calculate the probability of calculated stocks exceeding storage capacity or being negative by any given amount. We could therefore calculate "expected" overflows and shortfalls, and based the stock penalties on these.

Although there was an element of subjectivity in assigning penalties for high and low rates and for changes in rate and in deciding the relative weights for the different penalties, the system worked very well and was run every day to decide the rate to feed to the SPS model. There was a provision for "management override", which was occasionally taken up. However, with the benefit of hindsight we never discovered an occasion when we would have been "worse off" by following the model.

Perhaps the most interesting feature of Martin's method of solution is that since the uncertainties are incorporated in the stock penalty function, deterministic dynamic programming can be used, allowing quick and cheap computation.

The next application I will describe was also developed in Kuwait and concerned the long term development of the oilfields there. We had read in Martin Beale's book of the application of Separable Programming to oilfield development problems [5].

The essential difficulty is dealing with constraints of the form

production rate $\leq$ number of wells $\times$ well capacity.

In Martin's example problem, well capacity was assumed to be a function of cumulative production, so if a new set of constraints defining cumulative production for every time period is inserted, well capacity for any time period can be treated as a function of a single variable and a relatively simple representation can be used.

In our case [3], we believed that we could fit relatively simple models to predict the reservoir state variables of pressure and water influx (using the "Systems Method" of Rowan and Warren [6]) and from these and wellbore flow equations calculate well capacities, but we could not express them as functions of only one variable.

Martin Beale introduced us to "Non-linear Mathematical Programming", or more specifically "Conjugate Gradient Approximation Programming" [7].

The first step in his method was to generate a linear approximation to the real problem, and to put bounds on what he called the non-linear variables (that is those which if given fixed values would make the problem linear, in this case the production and injection rates for each reservoir in each time period).

He realized that, if this problem was solved, the reduced costs of those non-linear variables could be regarded as partial derivatives of the objective function with respect to these variables.

He could therefore use the techniques of unconstrained optimization (hillclimbing, line searches) to search for better solutions.

A system consisting of control programme, matrix generator, Linear Programming system and report writer was developed as in Figure 3.

In the days when the best computer we had access to was a UNIVAC 1108, solving a long succession of relatively large LP's took a very long time. On the whole we were glad not to have to run such a model to control our day-to-day operations, but we did use it as part of our annual planning cycle.

We found it all too easy to feed the model with inconsistent data (for example the reservoir behaviour equations had to be consistent with estimates of maximum recoverable oil), which led to some very interesting results. The model seized on opportunities to create new oil.

When we started work on the model we were interested in the problem of how far the fields could be developed before being allowed to decline to maximize economic benefit. However, in the early 70's the Government imposed production limits. The model proved a useful tool for investigating how best to achieve these limits.

Necessarily, the description of the reservoirs, wells, separation plant and pipeline systems were fairly "broad-brush". Therefore it was essential to do more detailed

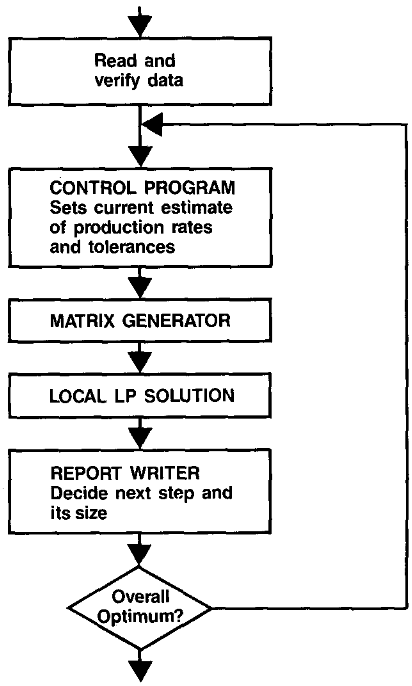

Fig. 3. Stages of the reservoir development model.

studies of the proposed solution before deciding on or implementing a plan, but at least we had a very good basis from which to start.

A further problem that arose was to find the right balance of investment in excess production capacity, storage capacity and loading facilities in order to meet predicted exports in a tanker fleet of changing composition (i.e. the advent of VLCCs): see Figure 4.

The problem was to minimize the investment in facilities, and the cost of tanker waiting time while keeping a reasonably constant production rate, so that gas supplies were assured and maintenance programmes could proceed.

We used a computer simulation model to evaluate various combinations of facilities. One of the things we discovered was that the "operating rules" were as important as the facility capacities [3].

In a simulation it really was all too easy, for example, to "build" extra storage tanks but effectively leave them unused. It would have been false to conclude in these circumstances that investment in storage was inefficient!

An entirely different application of computer simulation was used as a project management tool during the construction of the Sullom Voe Terminal in Shetland.

Critical Path Analysis has been used for many years, and although valuable can at times be misleading. If the actual time taken to complete an activity differs from that forecast, or the need for an extra activity arises, for instance remedial work if an item fails its acceptance test, or say repair of faulty welding arises, the critical

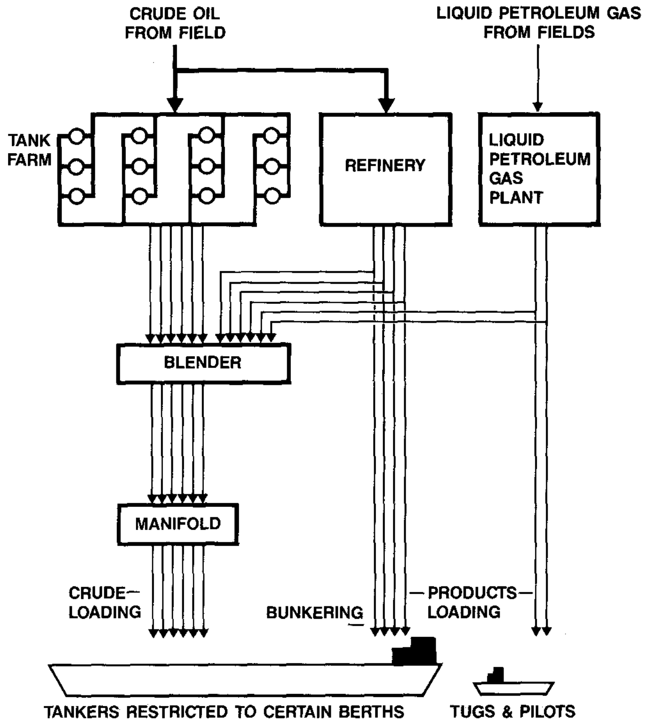

Fig. 4. Simplified tank farm and oil port installation.

path may change and management will have been paying too much attention to the wrong things.

A simulation model allows the variability of times, the effect of the weather, or the need for extra activities to be taken into account. Instead of a critical path, the output is the predicted "criticality" of each activity. That is the probability that an activity will become critical, based on its frequency in the simulated sample [8].

This was found to be a very useful tool, but if the models used are too detailed, it becomes very difficult to assemble and control the data, the computer runs are cumbersome and time consuming and the results do not become available in a timely and useful fashion.

The final group of applications I will describe concerns pipeline problems.

In the Kuwait Oil Company we had a network of pipelines through which oil was pumped from Gathering Centres to the two central tank farms. As part of a production expansion project, increased throughput from existing and new gathering centres had to be accommodated.

We had to decide whether to install new pipelines parallel to those existing on any arc of the network, or along new routes, and what the diameter of those required should be. Additional pumping capacity was also a possibility and we wished to standardize our pumps and prime movers as far as possible, so we thought in terms of standard units in series or parallel.

We could calculate pressure and flow distributions in pipeline networks, but because of the number of possible combinations this would have been a tedious business. Martin Beale, however, was able to suggest a formulation and solution using integer programming techniques.

There were non-linearities in the problem:

(a) The pressure versus flow relationships in the pipelines.

(b) The pump characteristic curves and net positive suction head requirements.

(c) Product terms with variables which took on the value 1 if a particular size pipe was built along an arc of the network, 0 otherwise. For convenience Martin called this the probability of the pipe being laid.

Flow along the arc was restricted to be less than or equal to the sum of the probability of each size of pipe multiplied by the flow rate through it; existing pipes were included, of course, with a probability of 1.

Martin's methods [9] of dealing with special ordered sets made these problems tractable and our problems were solved.

The methods used were generalised, expanded and used to study the possiblity of developing Gas Gathering Pipelines to collect gas produced in association with oil, and from gas fields, in the North Sea [10, 11]. The relationships represented were more complicated than for oil because gas is so much more compressible. The result of this exercise was negative, in that no such system was built, but at least it could be said that a very large population of possible solutions had been investigated.

A further development along these lines was used to investigate schemes for the redevelopment of the West Sole Gas Field in the Southern North Sea, shown in Figure 5.

The field had been in production for some years and, because of the drop in reservoir pressure, well capacity had declined.

The possibility of increased gas sales had arisen, so means were sought to revive production. Possible methods were:

(a) Drilling additional wells.

(b) Installing compressors offshore (which would have needed a new platform).

(c) Laying a pipeline (size to be determined) parallel to the existing one.

(d) Installing compressors onshore.

These options might be combined as in Figure 6.

The principal source of reservoir energy is gas expansion and so a relatively simple relationship could be used for reservoir pressure as a function of cumulative production. The drop in pressure as gas flows through the formation into the wellbore can be calculated as a function of flowrate. Similarly the pressure drop up the wellbore to the platform can be calculated as a function of flowrate. Using these relationships, well production rates and hence reservoir production rate in each time period could be treated as variables, so that the solutions gave a suggested production profile as well as a system configuration [12].

More detailed appraisal of the solution proved it was sound, a new contract was agreed with British Gas, and the project successfully implemented.

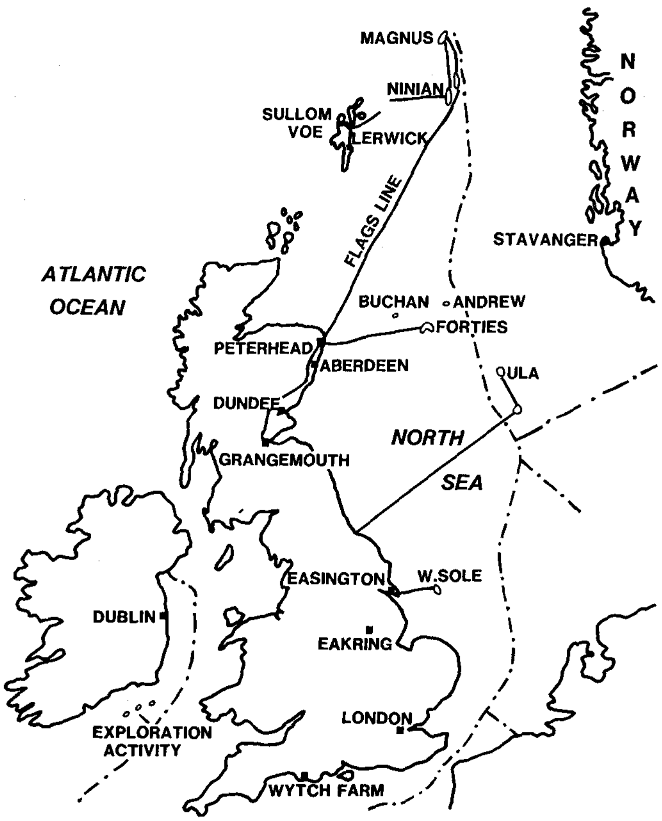

Fig. 5. Location of the West Sole gas field.

I said that I would concentrate my remarks on Exploration and Production activities, with which I am most familiar, but mention should be made of the very considerable amount of work done in integrated oil supply/refining/marketing problems using linear programming techniques; and also of very interesting work in the retail sales area, for example vehicle routing, depot location and filling station site selection and design, using logistics networks and financial and market models. Much of this latter work followed the intense competition for market share in the consumer based downstream market which arose when the major oil price increases in the 70's caused a drastic cutback in anticipated demand.

Comments on historical experience

There are many other examples that could be described, but I hope that the selection I have sketched has indicated the range of problems, tackled, the range of solution techniques used, and the usefulness of the results.

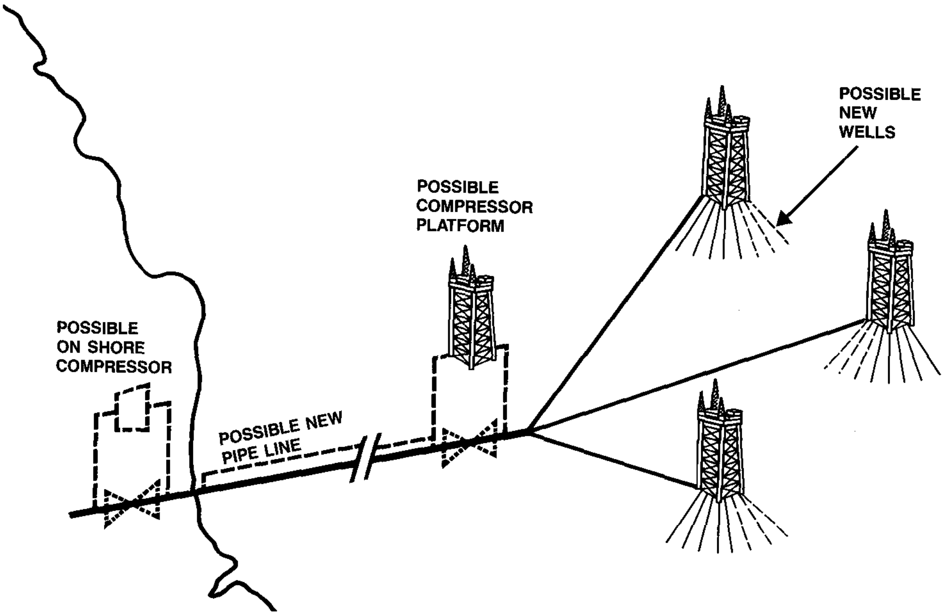

Fig. 6. West Sole gas field.

It has been said many times that a great deal of the benefit of using mathematical models lies in the greater understanding of and feeling for the problem that arises from the act of building the model. We certainly found this to be true.

Especially when using mathematical programming we found that "artificially" imposed constraints such as production targets and quality requirements could produce peculiar results, so often extra variables were inserted to allow these constraints to be relaxed at some notional cost.

It is not easy to avoid treating models anthropomorphically. Sometimes they do seem to take on a life of their own and take what one says all too literally—causing one to reply: "I didn't really mean it like that exactly".

Another feature of models is that necessarily they can deal with only part of the real world, described in a less than fully detailed way.

There are two consequences of this. Firstly the results depend, sometimes very heavily, on the assumptions made about those variables which are exogenous to the model, especially on assumptions about economic factors.

Secondly, since the models are approximations it is often necessary to do a great deal of further work before approving or embarking on a project based on model results.

The models described were in fact developed in the 60's and 70's when the world economy seemed to be stabler than it does now. We would need to review carefully what our objectives really are and consider the inclusion of extra constraints for protection against adverse outcomes in current unstable conditions. These models did not reflect our feeling that we should in some way "hedge our bets", and I am not sure if it is very easy to define exactly what we mean by this when we are dealing with a combination of political, economic, geological and technical risks over our whole portfolio of activities.

Current status

It is a fact that the current level of activity of the sort described, at least in my own company, is not very high and perhaps I could spend a little while pondering why this is so?

I think firstly that some of the ideas have been incorporated into the routine procedures and therefore we cannot always recognize Operational Research as a separate activity. For example, Risk Analysis is a routine feature of evaluation of Exploration Prospects, although problably more work needs to be done to see how effective this is, and how it can be improved.

I think that there are more important factors, however.

Management must spend most of its time on what it perceives to be the most important issues.

When one is faced with the problem of whether it will be technically feasible at all to develop an oil field in say deep water in Northern Latitudes, it seemed less important to fine tune the solution. Furthermore in the early days of development of oilfields, for instance in the North Sea, rapid development was regarded as the objective with less consideration of minimising costs whereas now the objective of minimizing costs is all important.

It is also true that the effects of Governments' licensing and fiscal policy have much more effect on our results than any other factor, so a great deal of effort is devoted to dealing with these problems. Also, for a variety of reasons we are involved in many more joint ventures and partnerships—the resulting blossoming of committees absorbs a great deal of time!

Yet another factor is that our technical "experts" have tended to become more specialized, and jealous of encroachments into their territory and their data. Evaluation begins by a geological definition of a reservoir's size, and goes through a reservoir engineering estimate of production mechanism and "capacity" versus time; with an engineering estimate of the structures, facilities, and costs to handle the prescribed production; and an economic analysis of the viability of the project in an uncertain future economy; this sequence tends to proceed with very little time for repeated cycles, and very little consideration of whether the project is the "best" use of the opportunities available.

Some of the "jealousy" has good cause; engineers for example find it very difficult to estimate for a vague or indefinite requirement and do not wish to be embarrased by an inappropriate use of "their" data in a world where a single offshore development costing £1 billion is not rare.

Also it seems that mainstream Operational Research has become a specialization with its own "high-tech" barriers. My feeling is that a set of solution techniques seeking problems is less productive than problems seeking solutions.

There have been changes in my company in organization and management style.

We have split our company into business streams and have a much more devolved style of managment.

The use of the large supply/refining/marketing models has declined; perhaps they were seen as a tool of centralized control. The early use of models to study the development of crude oil resources has not been followed up.

The reasons for this are probably that the role of the model in communication and integration has not been sufficiently recognized, or has been allowed to decline so that model results came to be seen as externally imposed. If communication does break down the models are not kept up to date and perhaps they end up solving yesterday's problems! There is also a tendency in these circumstances for the "modellers" to spend their efforts on the solution techniques and the efficiency of the number crunching rather than on the problem to be solved.

The future

As I have indicated we have had considerable success in the past in using OR to the benefit of our business, but these successes are not always being followed up.

But whatever our management style or organization, there is still the need to try to ensure that all parts contribute to the good of the whole.

If our problems have become more difficult, because of increased technical difficulty, because of more difficult economic conditions, or because of greater uncertainties of all kinds, it is less rather than more likely that we will produce good solutions by "guessing".

Therefore we are missing an opportunity if we do not use the very powerful tools which OR has made available, or the philosophy which is its basis.

Modern information technology provides very much more powerful means of data storage, computing and communication than in the past, and there has been enormous development in software. The tools are available for a revival of OR.

What else is needed to bring it about?

Perhaps we should examine our organization to see if it is too compartmentalized. We need to educate management and technical staff, convincing them that models provide a means of integrating their efforts and enabling them to communicate. We need to persuade them particularly that this communication is about relevant matters. As a start we should integrate our corporate data.

We need to educate OR specialists, so that they do not concentrate too much on the algorithms, but more on the problems and on their role as integrators.

Perhaps the easiest way to summarize this is to say that we should try to follow the example of Martin Beale, who, while making enormous contributions to theory, was at his very best when working on "real" problems.

Acknowledgements

I should like to acknowledge that most of the projects described in this paper were performed in Kuwait for the Kuwait Oil Company now part of the Kuwait Petroleum Corporation.

I should also like to acknowledge the contribution made to the preparation of this paper by colleagues and ex-colleagues in BP; in particular Jack Beasley, Graham Cuddeford and Keith Goodwin. All of them knew and worked with Martin Beale, and without them this paper could not have been prepared.

References

[1] E.M.L. Beale and G.C. Cuddeford, "The model-first view of OR" (1975), unpublished paper.
[2] B. Magee, *Popper* (Fontana Modern Masters, London, 1973).
[3] H.M. Ali, J.F. Beasley, A.S.J. Batchelor and E.M.L. Beale, "Mathematical models to help manage the oil resources of Kuwait" SPE 6711 (1977).
[4] G.M. Cain and M.T. Shehata, "A computerised system to optimise daily oil and gas production in Kuwait" SPE 9622 (1981).
[5] E.M.L. Beale *Mathematical Programming in Practice* (Pitmans, London, 1968).
[6] G.R. Rowan and J.E. Warren, "A systems aproach to reservoir engineering: optimum development planning," *Journal of Canadian Petroleum Technology* (1967) 84-94.
[7] A.S.J. Batchelor and E.M.L. Beale, "A revised method of conjugate gradient approximation programming," in: A. Prekopa, ed., *Survey of Mathematical Programming*, Proceedings of the 9th International Mathematical Programming Symposium held in Budapest, August 23-27, 1976 (Akademie Kiado, Budapest, 1980) 329-346.
[8] A.M. Feiler, "Project risk management," delivered at Internet 72 Congress, Stockholm, May 17, 1972.
[9] E.M.L. Beale and J.J.H. Forrest, "Global optimisation using special ordered sets," *Mathematical Programming Study* 10 (1976) 52-69.
[10] E.M.L. Beale, "The optimisation of gas gathering pipeline networks," *Bulletin of The Institute of Mathematics and its Applications*, May/June 1979, Volume 15.
[11] Department of Energy, "Gas gathering pipeline systems in the North-Sea," Energy Paper No. 30 (HMSO, 1978).
[12] E.M.L. Beale, "A mathematical programming model for the long-term development of an offshore gas field," *Discrete Applied Mathematics* 5 (1983) 1-9.