The Role of Interactive Colour Graphics in Boosting User Understanding and Confidence

Dr I McIvor and Dr D E Probert
British Telecom
Long Range and Strategic Studies Division
88, Hills Road, CAMBRIDGE,
England, CB2 1PE.

ABSTRACT

Our role is to advise senior British Telecom management on strategy for BT as a whole. This requires coherent strategic analysis aided by systems dynamics models. All management levels must have confidence in the models and their results. For analysing alternative futures we find that graphs are easier to appreciate and understand. We have also found that colour graphics greatly enhances the presentation of more than one curve at a time. Interaction with models in real time is a major step in boosting user confidence for it allows rapid confirmation, or rejection, of the user's prejudices.

Interfaces to computer models, such as menus, bit pads, etc, are successful if they interface efficiently between the user's mental map and that which is enshrined in the model. If the user can move a 'lever' which exists in the real world and that causes the model to display the effects he expects, then he will have confidence in the model. Decision makers want the best strategy. We shall discuss how we use colour graphics to compare strategies, but that often begs the question 'Why?'. This requires techniques used in artificial intelligence, which can also be used to 'customise' interfaces to individual users.

1 INTRODUCTION

The Long Range and Strategic Studies Division of British Telecom has the role of advising policy makers within the business. In recent years, technological and political developments have been having an increasing impact on the operations of telecommunication administrations. The effects of these are felt in all areas of business operation within British Telecom and require a response from the business as a whole. This response must be coordinated by a well defined strategy, the development of which requires coherent strategic analysis. To help in this analysis systems dynamics models have been produced by the division to cover many levels of the business from global views of finance to details of networks or terminal markets.

The development of these models is a topic in itself and will be discussed briefly in order to introduce our involvement in colour graphics. However, the interaction and dialogue between user and model also requires design, and we take this domain as the main theme for this short paper. A model can be very good but if nobody understands it or is willing to use it then it may as well not exist! The criteria users apply to a model before deciding to use it, or its output, have been changing with time. When the first model was written it was used only by its author, who fully understood its output and internal reasoning. Since then more people have begun to use the model and their requirements have been from a more naive viewpoint. We shall discuss these changing requirements and the methods used to satisfy them. The methods used are, in hindsight, quite simple and obvious, but are not regarded as being sufficient. Finally, we shall discuss the next step which we feel is at least as important as all of the previous developments combined - and probably more revolutionary.

2 THE MODELS

In 1976 a project was started to develop the Long Range Planning Model (LRPM) [1] with the help of the Cambridge University Engineering Department, and has been in use since 1978. British Telecom had been using forecasting models prior to 1976, but this was the first model which could be said to be truly strategic in its orientation. The main aim of the model was to be able to treat the organisation as an integrated system, and not a number of separate departments. In this it has been successful. It consists of four main interactive modules to span finance, the marketplace, personnel and telecommunications technology.

In 1979 a strategic control unit [2] was added to the LRPM in order to track corporate objectives and to allow the study of crises. The method used was not itself complex, though the definition of objectives and the kind of 'decisions' the control unit could take were not easily understood by the naive user. This unit is not now used as much as had been anticipated, for reasons we shall discuss below. Crises are still studied with

the model, but the user normally makes the decisions himself and then feeds them into the model.

By 1980 it was evident that telecommunications market in Britain was going to be liberalised, and that the LRPM was not able to cope with such developments within its own demand module to the depth required by many parts of the business. For example the model considers demand as falling into one of three categories, telephony, data and wideband communications. At a corporate level this was sufficient, but not within the service departments. To cope with this a new model was designed, the Integrated Communications Demand Model[3]. This was constructed with the knowledge that new services, terminals and networks would be introduced over the years and it was designed to be able to expand to cover this change. This has also been of use, but it is restricted by the same problems of understanding and maintenance as beset the control module of the LRPM.

The LRPM is, by necessity, a very aggregate model and the Integrated Communications Demand Model is much less aggregate. However, it is still not able to look at the marketplace in anything like the depth that is expected of a forecasting model. It is relatively large, being several thousand lines of code in length, even ignoring the internal documentation, and making changes to it is not as easy as one might have wished. Also there had been a large amount of discussion about networks and the ways that they might evolve and be cross-connected. This led to the development of a Network Model which enables one to explore the plethora of network alternatives which now confront the customer for telecommunications services. It is intended that once the model is mature it will be used to replace the network module within the Integrated Communications Demand Model.

To summarise, we have three models in use from a fairly simple Network Model up to a very aggregate, complicated Long Range Planning Model. Initially each of these models was used solely by its designer, but now none of them is under the designer's direct control. Indeed, in the case of one of the models it is even difficult to contact the original designer, but this has still not prevented us from exploiting it to its full potential!

3 USE OF THE MODELS

The models were designed by a research team in the Department of Management Systems, Cambridge University working under contract for British Telecom. There were thus significant hurdles to be overcome before the models could be implemented as a form of in-house decision support for senior BT management.

We see three key problems which confront the current users:-

1. Understanding the models sufficiently to use and aid them

in making their recommendations to senior management.

2. Convincing senior managerial staff that the models are sufficiently well modelled on reality that they are useful.

3. Convincing senior staff to use the models either directly or through passing strategic issues down to the modelling specialists.

The first problem includes the more basic one that the user should be able to access the model and use it in a reasonable time. In order to understand the model the user must be aquainted with such details as its internal logic, assumptions, and conventions, to a level sufficient to enable him to map it onto his own internal understanding of the world. In particular, the user must understand the input policy parameters, the output indicators, and the relationship between the two.

Senior staff require well presented reports if they are to be able to use them properly. Sometimes this may be possible using text alone, but the contents of graphical output are much more readily absorbed. Even so, a graph must be backed up by reasoned argument, or at least a belief that the person who produced it knows what he is doing and can be relied upon. If that person used a model then the model must also be believed by senior staff. Their understanding need not be anything like as detailed as that of the modelling specialists, but there must be some knowledge pertaining to the underlying assumptions and structure.

The models are never used in the absence of background information, but to check the ideas and intuition of the user, or group of users. Similarly, the results obtained from the models are not presented in their raw state, but as part of a reasoned analysis. In the group context the model becomes almost like another member of the group. If the model is internally consistent then it is the only "member of the group" that can be relied upon always to give the same answer to the same question - assuming no change to the model or its other inputs. This is both an asset and a restriction, as we shall see below.

4 THE INTERFACE TO THE MODELS

In one respect the models are the same, they use a common interface so a user does not have the problem of learning a new interface when moving from one model to another. The models are all run on a mainframe over the public switched telephone network. This introduces the problems of access to a well used machine and the limitations of speed implied by an ordinary telephone line. These are currently part of our environment and so we shall ignore them even though speed can be very important.

Initially the Long Range Planning Model was accessed by a slow teletype. This was a cumbersome terminal to use for an interface. It required the user to remember a variety of commands, was slow, could not produce acceptable line graphs and certainly put off all but the most determined users! The model was able to produce graphs offline, but only on a line printer as a graph plotter was not available. Line printer graphs are difficult to read, even for a single curve, and are confusing when several graphs are output. Further, as the line printer was a day or more away by post, one had always to produce all the graphs one needed without knowing if they would be giving the required answers. Of course the analyst could always list the graphical output at the terminal, but a complete set of graphs even for just one run would take several hours on the slow, 300 baud teletype terminals. So in summary, this original user interface for systems dynamic modelling produced plenty of scrap paper, in which the really useful results were still buried in a veritable flood of information.

The use of a higher speed visual display unit helps the user, both in terms of response speed and the ability to look quickly at a graph before a hardcopy is taken. However, the graphs are still of a line printer form and suffer from the same defects of readability and confusion. Before that sort of graph can ever be presented to senior management, it must be redrawn by hand on proper graph paper - a tiresome job. For many applications the accuracy implied by line printer graphics is insufficient, but in our use such accuracy is not required since we do not claim to forecast. In fact such accuracy is to be avoided in order not to give the impression that we are producing rigid deterministic predictions.

Therefore it was decided that a graphical terminal was required. There was the need to be able to plot many graphs on one screen without confusion. This meant either a very complicated scheme of dot-dash patterns to distinguish between curves, or the use of colour. The first had to be used for hard copies as the cost of colour photo-copying was too high. However, for on-line working on the terminal the use of colour was seen to be the obvious solution. It enabled the user to see at a glance which curve was which, and further allowed presentation to senior management on-line. This latter point was regarded as the first step towards being able to convince non-users of the reliability of the model without the need for a long process of convincing them by the correctness of results. Non-users were able to ask for a particular output to be displayed to ensure that it coincided with their ideas. Examples of output are given in Figures 1 and 2.

We saw four areas in which a colour graphics interface was required:-

1) Design of a model. The interface was not available when the LRPM was designed and built, but was for the other

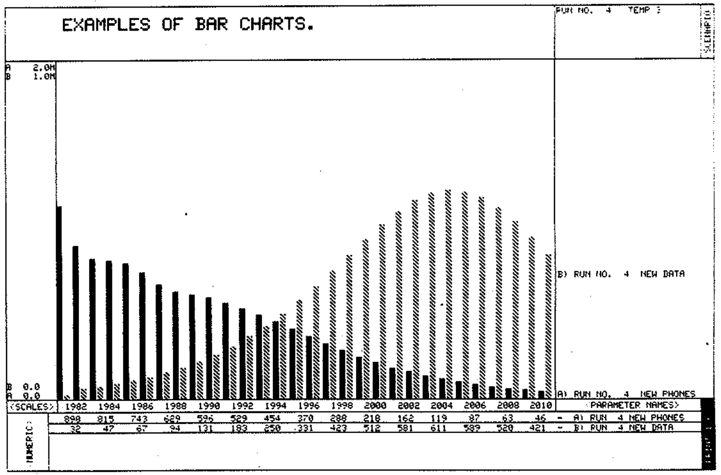

These results are merely examples and do not represent BT policy.

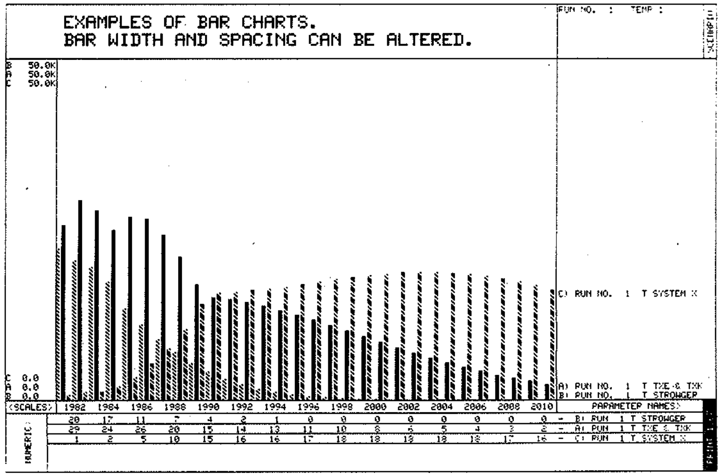

Figure 1

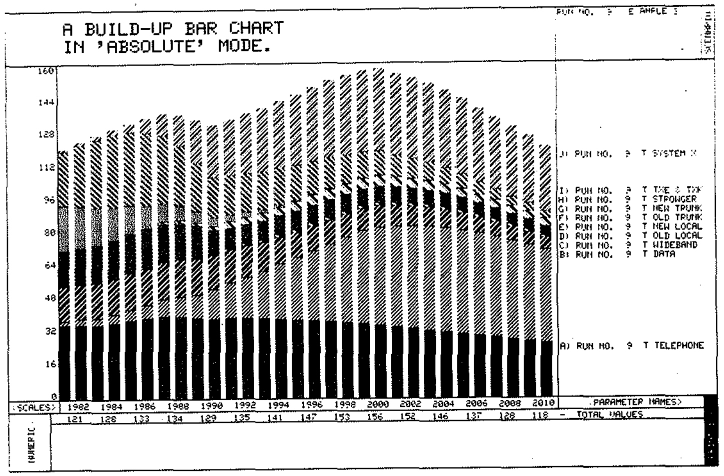

These results are merely examples and do <u>not</u> represent BT policy.

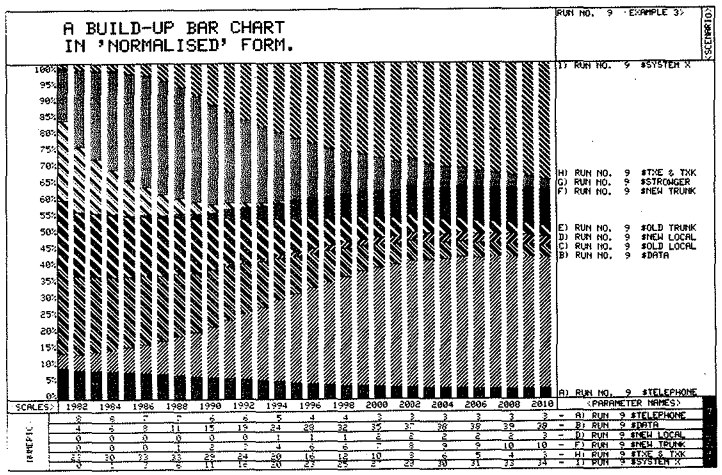

Figure 2

models. The advantages, in terms of both time and effort, were considerable. Not only did it help that we had a basic framework into which the new models could slot, but the ability to plot graphs after even the smallest of changes to the code was an aid to speedy development by promoting the designer's confidence in the model being implemented.

2) Validation for new users. This is covered in the following section on user confidence.

3) Use of the model. We need to be able to exploit our models to the full and the interface helps us to do that in ways discussed below.

4) Maintenace and tuning of the model. As with the design of the model the availability of the interface helps tremendously. Minute changes of the policy parameters when tuning to a new base view of the future can be evaluated and their effects seen rapidly. Changes to data held in the database can also be made and checked rapidly. Finally changes to the code of the models can be checked for consistency and for their ability to match their specification.

The move to colour for interactive graphical output [4] helped considerably, but in itself it was not enough. Further additions were made in the form of a menu on the screen and a bit pad. This enabled users to see the options open to them at any point, and so concentrate their attention. The form of the menu was colour coded boxes with labels, and the user could move cross hairs to the particular box required. The use of a bit pad means that all possible commands are on display at all times. The commands are structured on the bit pad so that all the graphical display commands are together, etc. As with most advances, this was felt to be a significant step forward at the time, but in retrospect it is probably not of major importance. Neither type of input has really brought out the structure or logic within the model, though the bit pad has helped to bring out the structure of the interface, which is of use to all users, naive or experienced. Attempts to bring out some of the logic have been made by grouping together policy parameters in meaningful groups - but this still does not indicate their possible effects. The colour graphics workstation is shown diagrammatically in Figure 3 and its important place in the user interface is shown in Figure 4.

A recent addition to the software helps a little, though again the logic is not brought out. A facility was developed which enables the user to see a colour coded comparison of the results of one run against another, or base, run. It measures the length of the difference curve as a percentage of the values in the base run. The idea is that an output indicator which has a large variation in the second run compared to its values in the base run would show up red, with a movement through the spectrum

COLOUR GRAPHICS WORK STATION

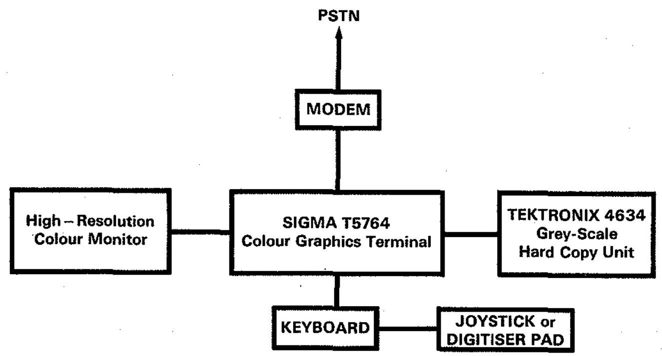

Figure 3

THE USER INTERFACE FOR STRATEGIC MODELLING

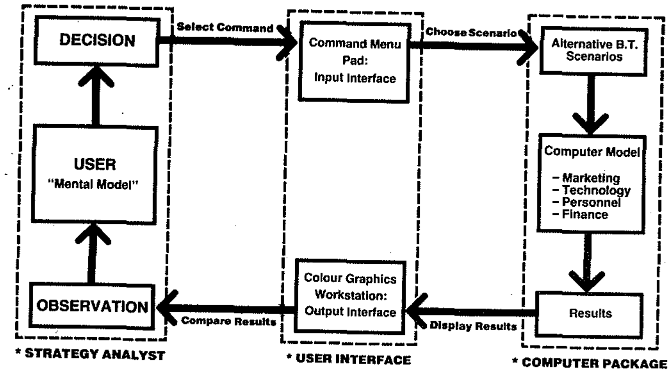

Figure 4

as the variation becomes smaller. It is particularly sensitive to oscillations and less so to constant deviations. This has certainly helped users to determine which indicators to investigate first, or to show up changes which were not expected. There has been discussion about the definition of the measure to be used but the agreed problem is that it still does not help the user decide why the indicator changed or what to do about it.

Within the system there is a videotex database of help pages which use the colour graphics. These were originally designed to help users when they wished for explanation of the model, the interface, the output indicators, etc. It is not currently used for this purpose, but for giving demonstrations of the system to potential users of the models in which the displayed information is still of a more general nature.

5 USER CONFIDENCE

We hope the above has conveyed the message that interactive colour graphics helps in gaining the confidence of both users and recipients of material based on dynamic computer models, but that there is still a long way to go. Some people may argue that colour graphics tends to cloud the issue, in that it diverts attention from the results being displayed to the fancy capabilities of the terminal. For a while this is certainly true - but after only five minutes of argument about a particular

problem, the colour is forgotten and any criticism usually focuses upon the scope of the output indicators themselves. The models usually have 144 output indicators, which is clearly not enough to answer <u>every</u> question a person may ask about every problem a model can address. There are three types of output which may be required, but which are not available to the user:-

1) Indicators which are used within the model, but not held as output indicators. These can be brought out by adding them to the list of output indicators through changing an output routine.

2) Indicators which are either aggregated or disaggregated within the model to a level different to that required by the user. Again this can be rectified by changing the output routine with the necessary change to the level of aggregation.

3) Indicators which the model does not yet consider. For example the LRPM did not have indicators of market share before we introduced new equations representing the dynamics of market competition. Such changes can involve many existing subroutines, and possibly the addition of new subroutines.

The LRPM has over 300 policy parameters which a user <u>can</u> change if required. This list is very daunting to new users who tend to think that they have to give values to all of these. In fact default values are given to them which give a reasonable, surprise free future as computed by the model. This future is generally agreed to be acceptable as a base projection and it is around this that discussion revolves. Thus a particular topic may involve the consideration of only a handful of policy parameters, the other 290+ being left unchanged. There is a second problem, however. The user must understand what is meant in his terms by a particular policy parameter. For example there are three parameters in the model, Total , Scale and Ascent for "Old to New L-Net: Technology transform for Local Network". One is able to understand a certain amount from that name but it is not easy - and, even worse, it is not obvious what one expects as a result of changing the values. To overcome both the problem of a large number of parameters and a lack of comprehensibility a system of indices was introduced. An index links together many individual policy parameters and allows the user to change them en masse by a percentage of their base values. The three parameters mentioned above are thus combined, with others, in a Technology Index which changes the rate at which new technology is introduced. This Technology Index is much easier to understand and naturally boosts user confidence in the model as it maps onto his own mental model of the way the world operates.

The user needs to have results explained, either by display of further results, or by the display of the logic or input. The logic used in the models could be stored in the videotex

database, as some of it is, but it would be in a very general form and not related to a particular question which might arise. The need is for the model to be able to explain why it came up with a particular answer. For example, in one evaluation using the LRPM the profit target was increased and the result was a long term reduction in tariffs. At first sight this is counter-intuitive, but in fact is quite logical as the higher profits mean an increased ability to pay off loans, hence less interest and eventually lower tariffs. In order to be able to explain this it must understand its own internal structure, something which none of our models can do at the present. The control module does give a certain amount of explanation in that the user can display the values assigned by the controller and the error vector, which is the difference between the values of the controlled indicators and the user's target values. However, even a sophisticated control module still does not provide "intelligent explanation" of its actions.

The current developments in "Expert Systems", or the more general Intelligent Knowledge Based Systems, seem to be the answer to at least some of these requirements [5]. A model using these techniques would not only know its own internal structure, but could keep track of the logic through which it goes and then report on parts of it if requested after doing an evaluation. An ideal system should converse in the user's language (terminology) and also be able to prompt and help the user. It should notify the user if, say, profits go negative and then suggest the user changes a particular policy parameter in order to rectify the situation.

Further, in the group context one expects an individual to change his attitude to a particular topic in the course of discussion. Normally this is due to the introduction of new information, the removal of some constraint the individual has applied, etc. The models at present can not learn in the same way. If one gives them the same policy values they will always return the same output indicator values, unless the code is changed. The only way in which one might describe them as changing due to introduction of ideas is when policy parameters not directly related to the discussion topic are changed. An "Expert Model" should have the ability to learn. For example it should be able to remember previous results in a way which will enable it to make some judgement about expected results before doing a complete run. It should be possible to inform the model that reduced profits over a short period are acceptable so long as over the long term they are enhanced. In such a system one would require a structure such as that shown in Figure 5. Here the interface knows what information is stored in the database and what the model, or models, can do. Hence, it can direct a particular query to the relevant information or model, or use both stored information and models to produce an answer - again storing any new information away in the database for future reference.

AN EXPERT SYSTEM FOR POLICY ANALYSIS

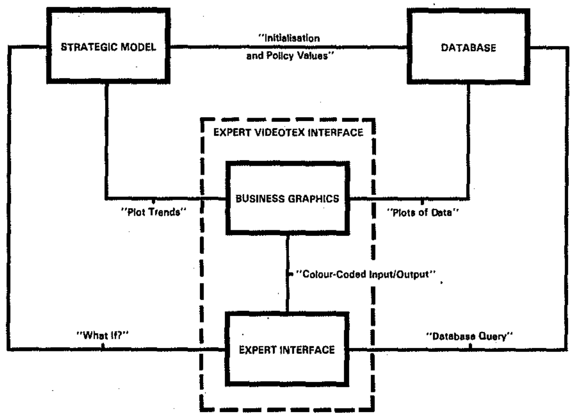

Figure 5.

6 CONCLUSIONS

In summary, we see the advances we have made to date being able to answer some of the questions users may pose. We see three main types of question which a user may pose after using our models:-

1) Where have changes occurred as a result of the new policy assumptions? These are answered by using the colour graphics capability, and it is the matching of these to the users intuition and prior knowledge which give the user confidence in the model.

2) How has the model dealt with the mismatch between the user's targets and its calculations? This question is answered by the control module with its output information. This again uses the colour graphics interface but gives the user confidence in the control module more than in the model as a whole.

3) Why has the model calculated the answers it has? This is still beyond the capability of the models, and requires the techniques of intelligent knowledge based systems. Once this capability exists, together with the ability to change the logic interactively, then users should be able to understand the models more completely and so gain yet another level of confidence in them.

So, while we have made significant steps, if not strides, in boosting user confidence by the use of colour graphics, we can see our next step forward as being the use of intelligent knowledge based systems. Such systems could be used in both the creation of "intelligent" user interfaces for system dynamics models as well as in the implementation of complete modelling packages.

7 REFERENCES

[1] Probert D.E., 1981, "The Development of a Long Range Planning Model for the British Telecommunications Business - From Initiation to Implementation", J. of Opl Res Soc (UK) Vol 32, August 1981, pp. 695-719.

[2] Probert D.E., 1982, "Systems Dynamics Modelling within the British Telecommunications Business", J. of System Dynamics, Winter Issue 1982. Presented at the 1981 International Conference on System Dynamics Modelling, Albany, New York, USA.

[3] Doubleday C.F., Probert D.E. & Walsham G., 1983, "A Strategic Model of Communications Demand", OMEGA Vol 11 No 4, July 1983.

[4] Allenstein B. & Probert D.E., "The Integration of Business Graphics and Videotex for Corporate Planning : A Case Study", State of the Art Report on "Business Graphics", Pergamon Infotech, June 1982.

[5] Probert D.E., 1983, "Towards Expert Systems for Telecommunications Policy Analysis", Computer Communications, April 1983. Presented at the British Computer Society Conference on Expert Systems, London, December 1981.

NOTE The views expressed in this paper are those of the authors and do not necessarily represent those of British Telecommunications as a whole.