7TH WCTR
SYDNEY AUSTRALIA
1995

TOPIC 13
PUBLIC SECTOR
PERFORMANCE

AUTOMATING ROUTE ALIGNMENTS: A NEW PHILOSOPHY

PETER G. GIPPS
CSIRO Division of Building Construction and Engineering
PO Box 56, Highett, Victoria 3190
AUSTRALIA

Abstract

Generating alignments which satisfy geometric and other constraints is a major branch of CAD but can take days or weeks depending on length, terrain and other constraints. Consequently, optimising alignments automatically drew considerable early interest, which gradually declined as promising approaches failed to be translated into solid applications. The paper revisits the problem and examines the goal.

INTRODUCTION

Constructing a new road or railway, or realigning an old one can be very expensive with costs depending on the alignment selected. Costs are increased by long structures, by large volumes of cut and fill, and by unbalanced cut and fill where the discrepancy has to be dumped or borrowed. The cheapest solution would be to follow the natural surface of the land, but this is frequently impossible as the alignment is subject to many constraints such as geometric standards and trying to avoid or negotiate features or areas en route in a particular fashion. The problem is exacerbated in urban or semi-urban environs where the construction of a by-pass or the reconstruction of an existing road is also influenced by the location of services, existing roads and buildings, and the financial, social and political costs of land resumption. These complexities make the task of producing an acceptable low cost alignment a difficult art.

Generating alignments which satisfy various geometric and other constraints is a major branch of computer aided design and is performed interactively with any of a large number of packages running on a wide range of computing hardware, from high-end PCs to mainframes. Depending on the length of the alignment, the complexity of the terrain and other constraints, this stage of the 'design' process can take days or weeks. As a consequence, the problem of optimising alignments automatically has been the subject of numerous efforts since the advent of CAD.

O'Brien and Bennett (1969) suggested a model of dynamic programming based on a rectangular grid to solve the problem of minimising costs. Nicholson et al. (1976) suggested a two stage model, in which the first (approximate) stage was based on O'Brien and Bennett, and the second on an unspecified calculus of variations. Parker (1977) suggested an alternative grid-based model for minimising construction costs subject to slope constraints. Trietsch (1987) developed a family of methods by which a solution may be sought which was based on four altemative grid searches and four different cost approximation methods. Although, Trietsch's analysis of the problems involved in computer aided route determination was very acute, the methods which he espoused did not really overcome them.

All of the methods involved some sort of search grid, and some method for calculating the costs for arcs between the straight sections of the route. However, the use of grid search methods meant that the number of points (and the size of the problem) grew very rapidly as the demand for accuracy in the location of the route increased. Further, when selecting a particular method the planner/designer had to decide whether it was more important to be able to make fine changes in direction or to make backward curves (backward curves occur when an alignment is so windy that the straight-line distance between a vehicle travelling the route and its destination actually increases). In many situations backward curves may be considered unnecessary and the choice is simple, but in mountainous terrain the use of backward curves and accurate placement of the route may both be essential to the development of low cost alignments.

Some more-recent authors (Goh and Chew 1988 and Chew et al. 1989) have attempted to place more emphasis on the calculus of variations which involves the construction of large systems of equations. However, these have a limited development path, as they increase in complexity very rapidly as more detail is required or new constraints are introduced.

None of the attempts to date have been sufficiently successful to be considered a standard design tool and be included in commercial packages and interest has gradually waned. Unlike CAD which has gone from strength to strength, the history of alignment optimisation is one of a massive early effort between 1965 and 1975 followed by a gradual abatement as promising approaches failed to produce methods that could be translated into solid applications.

From a current perspective, a major cause of the lack of success was the limited capability of computer hardware available when interest in alignment optimisation was at its peak. A modem well equipped microcomputer is close to 100 times as fast as a typical mainframe of 1970 and can have 100 times as much memory. This comparative lack of computing power (by current standards) limited the algorithms which could be used, and thence restricted the complexity of the problems attempted. While modern computers allow much more power to be employed, the application of that power via (modifications of) old algorithms is not necessarily the most effective approach. The availability of greater computing power means that it is important to address some basic questions; not just 'What algorithms are most suitable?', but more fundamentally 'What is the problem?'.

THE GOAL OF ALIGNMENT OPTIMISATION

Classically planning has a number of key questions which are repeated on a cyclical basis until the answers remain unchanged in successive cycles:

(a) Is the facility necessary?

(b) How big should it be?

(c) Where should it be located?

(d) What will it cost?

In relation to a road the questions translate into:

(a) Is a road connecting towns A and B necessary?

(b) What capacity should it have?

(c) What route should the road take?

(d) What will it cost?

Then because accurate information about costs could have modified earlier decisions, questions (a), (b) and (c) are repeated with cost information included. Modifications of the earlier decisions require a full or partial redesign and another cycle through the questions. The grosser limits on the location of the road may be determined by planners on the basis of environmental or social impacts. But the finer placement is left to the designers because costs depend critically on the placement, and only the designers have the tools to answer the question of cost. In practice the same people frequently fill the roles of both planners and designers, but because the need to be able to answer the cost question their training and orientation tends toward design. Thus in highway engineering questions (c) and (d) tend to become fused and passed over to designers to answer. This influences perceptions about where the boundary between planning and design lies and affects the goal towards which alignment optimisation is directed.

The first problem encountered in alignment optimisation is deciding the goal. This is not trivial and passing over it too rapidly can have serious consequences during program development. Deciding the goal requires that a number of issues have to be addressed. The first issue to be addressed is whether the alignment is to be a plan or a design.

Planning versus design

In this context planning is taken to include determining where the route should go and obtaining preliminary estimates of the cost; while design is concerned with defining the route, specifying earthworks and other construction details and producing more accurate cost estimates.

Worthwhile discussions concerning planning and design can only be undertaken if it is understood that there is a substantial subjective element involved in both. This is more obvious in planning where the optimal alignment represents a balance between low cost and non-quantitative factors relating to how the road impacts on the environment, users and the rest of the community. Such a balance by its very nature is subjective. But design also contains its share of subjectivity, though this is more implicit than explicit. The 'obvious' objective goal of 'minimising costs' does not mean minimising costs in the mathematical sense, but rather achieving a balance between low cost and meeting criteria that are generally accepted as constituting good design. For example, an increase in cost produced by eliminating a bend (which satisfies all the geometric criteria) may be acceptable if the increase in cost is 'not excessive'(in the opinion of the designer).

Bearing in mind this subjectivity, should alignment optimisation be part of the planning process or part of design? The planner is concerned with locating the alignment in such a way that an optimal balance is obtained between costs and adverse impacts, and is not really interested in the details of the design. The designer is concerned with the final definition of the alignment and precise costs; and the planning role of route location was obtained by historical accident.

The difference between planning and design affects the way in which details are viewed; that which is important to a designer may be irrelevant to a planner. Thus an alignment which may be good to a planner may be deficient to a designer. For instance, one criterion of good design is that nearby curves in the same direction separated by a short straight should be replaced, if possible, by a single curve based on a large radius arc. This criterion is related to how the road is defined and if the refinement is feasible usually does not affect the location of the road greatly. Thus the difference between a good planning alignment and a good design alignment lies in the definition of the alignment, not in the location of the road.

Depending on the stage which planning has reached, locating the alignment in plan view to within a hundred metres or five metres may be sufficient to enable the impact of the facility on the environment and the community to be assessed. However, a major consideration in planning is value for money, and costs can only be reliably estimated once the alignment has been adjusted to minimise earthworks and other construction costs. This is a difficult process, and currently requires the skills and computational facilities of the designer. Thus effective planning requires design not for its own sake, but to provide the numbers that are essential to assessing and revising plans. One consequence of this dependence is that a lot of planning is misclassified as design.

An optimal design?

Is there any such thing as an optimal design? It is quite legitimate for the senior designer in an organisation to claim that a particular alignment between A and B is the optimal design given the terrain and various constraints. However, it is not possible to pass from the specific to the general. There are differences between the practices adopted by different organisations that mean a design accepted by one as being the best they can achieve in the circumstances, may be regarded by another as requiring substantial revision. Even within an organisation subjectivity can produce substantial differences between designers.

The problems of differing design practices between organisations and differences between opinions of individuals within those organisations, mean that it is not feasible to contemplate constructing an optimal design automatically. The concept of an optimal design is too elusive in reality to contemplate automating it. Rather attention should focus on the planning side; producing an alignment which is good from the planning perspective and letting the designers amend it to meet their criteria.

An optimal plan?

Is there any such thing as an optimal plan? Again subjectivity raises its head, but now many of the issues are contentious and political. What is the value that can be placed on ensuring that a freeway is kept out of sight of a scenic lookout? What is the social cost land resumption? The cost of an alignment intruding into certain areas, such as private properties, is a non-linear function of the extent of the intrusion as extraneous (and frequently unquantifiable) costs such as legal charges and public relations have to be considered. Differences in construction costs between an alignment that intrudes and one that does not have to be balanced against the costs of acquisition, public relations et cetera to determine whether intrusion into the zone is worth considering.

The question of subjectivity and the different approaches between organisations and individuals make the question of whether an optimal plan exists somewhat problematical. In fact, an optimal plan in general is probably even more unlikely that an optimal design in general. The question now arises: 'If there is no such thing as an optimal design or an optimal plan is there any future to trying to optimise alignments automatically?' The answer to this is 'yes'. The fact that organisations actually plan and design roads interactively means that optimal plans and designs do exist in the particular if not the general.

The alternative to trying to find an optimal plan in general is to recognise that there is a multiplicity of locally optimal alignments that have differing qualitative properties; and instead of a trying to find a single global optimum, produce a selection of relatively low cost alignments which meet the objective criteria. The planner is then free to select the alignment which represents the optimal balance between low cost and minimal adverse impacts (in his opinion) and to amend it if necessary.

CREATING REALISTIC MODELS

There are many factors which need to be considered during planning and design and which can have a major impact on the optimum alignment of a route. Unless a model can cope with a minimum range of factors it cannot provide the tools necessary to produce realistic solutions. To provide a minimum level of flexibility a models needs to incorporate information about:

Geometry

The basic description of the geometric standards of the alignment must include minimum radii of curvature, maximum gradients, and location, bearing and gradients at the end points.

Terrain

The shape of the terrain affects the volume of earthworks.

Geology

The costs of earthworks and structures depend upon the geology. A transport corridor of necessity must contain at least one geological zone. Further, each zone can have a number of strata, with their own individual characteristics of batter, shoulder width and extraction costs.

Linear features

Most corridors will include linear features such as roads, rivers, railways and pipelines which may have to be crossed. Some crossings must be at the same level as the existing feature, others must be at a different level and may or may not involve a structure. The minimum description required of each feature includes the location of the feature, the nature of the crossing, and horizontal and vertical clearances.

Special treatment zones

In addition to linear features there are often zones which require special treatment for social or environmental reasons. A railway line may need to avoid a park, or a road a water catchment. Even when the alignment can pass through a zone it may be forced into a cutting or tunnel to reduce noise or visual impacts, or onto a viaduct with a minimum elevation when crossing a flood plain.

HANDLING SUBJECTIVE FACTORS

In order to provide a satisfactory alternative to current methods, a user must be able to address subjective issues such as environmental impact and land resumption. The subjective factors involved in generating an alignment can be extremely complex. As an illustration consider a simple example in which the alignment may intrude into a National Park which has been established to protect the flora and fauna native to a certain river valley. The Park includes the entire water catchment at the head of the river, and extends over the ridge on either side of the valley to exclude exotic weeds and other water carried pollution. This region outside the valley is part of Park but is not part of the environment that the Park was established to protect; rather it is a buffer zone which helps protect the environment of the valley. Because of the way in which land for the park was acquired or the boundaries defined when the Park was Gazetted, this buffer zone is not of uniform width. As a consequence, intrusion into the buffer zone can be of varying importance:

- Where the buffer zone is wide intrusion of a road or railway into the outer parts of the buffer zone may have negligible impact on the environment of the valley even though the area of the intrusion may be large.
- Where the buffer zone is narrow or the alignment penetrates close to the watershed of the ridges dêfining the valley, the efficacy of the buffer zone can be threatened even if the valley itself is not entered.

Penetration of the valley by the alignment raises further possibilities:

- If the alignment crosses the valley at the downstream end of the park the intrusion, while undesirable, may only have a very local impact.
- If the alignment crosses the valley in the upper headwaters it may provide a means by which exotic weeds and pollutants can enter the system but otherwise leave the park relatively intact.
- If the alignment crosses the valley in its central reaches it renders only part of the valley open to infestation by exotic weeds and pollutants, but may provide a barrier to certain fauna that renders their habitat unviable.
- If the alignment runs along the valley it renders the whole valley open to infestation by exotic weeds and pollutants, provides easy access to feral cats, dogs etcetera.

Minimising objective costs while preserving environmental amenity is very complex as the questions of how environmental impact is measured and how it is balanced against objective costs are basically political, in the sense of being a matter of policy. There are two basic strategies which can be followed when attempting to cope with this situation. The first is to try and quantify the subjective aspects and to incorporate the resulting values in the objective function used in the optimisation. The second is to produce a number of alternative alignments based on optimising the objective costs, and then select between them on subjective grounds.

The first strategy requires that monetary values be assigned to the areas of environmental importance which are then incorporated in the objective function. This approach is not as simple as it might appear initially, as assigning a value to these environmental costs is extremely difficult. It may be possible (though difficult and costly) to allocate a monetary value to each square metre of the Park based on the impact that the alienation of that square metre would have on the integrity of the park, and to create a contour map of the environmental costs of alienation of individual square metres. But, this does not address the more difficult problem of connectivity, which considers the impact of the alignment on the contiguity of the Park. This approach is subject to the problem that any interested party that disagrees with the result will contest the values assigned to the subjective factors; arguing for an increase or decrease depending on how they want the answer changed. Further, it raises the philosophical problem of whether a computer program should make the decisions about subjective issues in this fashion.

The second strategy recognises that once subjective factors are introduced the problem becomes more complex with more than one solution. Environmental impact is essentially multi-dimensional so that any judgement about the value of environmental amenity also involves judgements about the value of different aspects of the environment relative to each other. The approach is to determine whether a conflict exists, and if it does what additional objective costs are imposed on the alignment by avoiding the environmentally sensitive area completely. If this is considered excessive, various compromises can be investigated. This is the favoured strategy and it is based on the tenet that computer packages should not make decisions that are essentially subjective, rather they should be used to generate objective information which can be used by planners as a basis for their subjective judgements.

General strategy

The general approach adopted in the strategy is first to identify whether conflicts actually exist by running the optimisation without including any information about the location of environmentally sensitive zones. There are three possible outcomes from this:

- The lowest cost alignments avoid such zones and there is no conflict.
- The lowest cost alignment runs through environmentally sensitive zones, but other alignments whose costs are not too dissimilar avoid them and are acceptable alternatives.
- All the lowest cost alignments run through environmentally sensitive zones, and no economically acceptable alternatives were generated.

Only in the last case is there a serious conflict, However, the fact that no acceptable alternatives were generated in the first stage does not mean that none exist. Their non-appearance in the first round of optimisations may merely be a result of the boundaries of the zone(s) occurring in regions where the objective costs can be reduced by moving the alignment into the zone(s). The next step is to investigate this possibility more fully by defining the environmentally sensitive zones as zones that must be avoided and reoptimising.

If the alignments obtained by constraining the alignments to avoid the sensitive areas contain an acceptable alternative the process can stop, otherwise it its necessary to seek a compromise. The compromise may involve allowing the alignment access to the outer area of the buffer zone described earlier, or allowing the alignment access to the core area of the park subject to certain constraints on the nature of construction.

OPTIMISATION TECHNIQUE

The generation of sets of low cost alternatives can be readily achieved by employing one of a number of stochastic optimisation techniques. A randomly selected intersection point defining the alignment is moved at random subject to chosen constraints. The cost of the modified alignment is calculated and compared with the cost of the old alignment, and a decision is made to accept or reject the change. A move is accepted if the new cost is less than the old cost plus a rejection threshold. Once a decision has been made, another intersection point is selected at random and the process repeated.

The rejection threshold is a device to enable the alignment to escape from local sub-optima. It provides the mechanism by which the alignment can accept changes that lead to cost increases on a temporary basis. Initially the rejection threshold is very high but as the optimisation proceeds the rejection threshold is slowly reduced to zero and finally only those moves which lead to a reduction in cost are accepted.

The algorithm does not guarantee that the alignment will escape from local sub-optima; it merely makes it possible. However, careful selection of the initial rejection threshold and the rate at which it is reduced can greatly increase the probability of reaching the global optimum. If the initial value is too low or reduced too fast, the alignment is more likely to be trapped in a sub-optimum. On the other hand, if the initial value is too high or reduced too slowly, the time required may be excessive and only a limited number of alternatives will be produced.

The stochastic nature of the technique means that the solution obtained from a particular optimisation run is not a deterministic function of the starting solution. Indeed if the initial rejection threshold and the rate at which it is reduced are chosen carefully the final solution is virtually independent of the starting alignment. Consequently, the resulting alignments are not constrained by preconceptions of where the route should lie. A process somewhat akin to lateral thinking. This isolation of the final solution from initial alignment is important because it removes one more element of subjectivity from the solution. Even though an expert can produce a far cheaper alignment that a non-expert, the cost of the alignment produced by the expert is critically dependent on his/her initial judgement of where the road should go based on impressions of the macro topography, while final costs are not determined until much later when the alignment has been adjusted to fit the micro topography.

The complexity of the land surface and constraints, and nature of the process, mean that a stochastic optimisation package does not yield the same answer each time. Rather, twenty runs are likely to produce twenty different alignments. The multiple solutions tend to be clustered in families based on their behaviour with respect to the more significant features of the corridor. Because the stopping criterion halts the optimisation when range of movement of individual intersection points is small but not infinitesimal, the alignments output are all approximations to the local optimum, and a particular local optimum may be represented by several approximations. Variations between approximations to a particular optimum are much smaller than variations between different optima.

This multiplicity of solutions has the very positive effect of answering some of the 'What if' questions that arise as part of the normal planning process. (What if we tried this variation? Or that variation?). If a variation is at all comparable in cost with the best solutions it will probably appear in the list of possible solutions. If a variation is relatively expensive it will be unlikely to appear. This allows the planner to make subjective judgements about the relative desirability of various 'low cost' solutions and balance these against the objective estimates of their costs.

BENEFITS OF APPROACH

The use of random moves for intersection points subject only to the geometric constraints under the stochastic approach leads to a number of major benefits.

Robustness and extensibility

Selecting a random location when shifting an intersection point obviates the need for calculating anything more complex than the change in cost produced by the shift. This ensures that the model is robust and readily extensible. That is, that it will work with a large range of diverse data sets and be able to incorporate new types of constraints without major rewrites.

The twin attributes of robustness and extensibility can be promoted by using only a few absolute constraints to define the permissible space for the new location of an intersection point, and manifesting other constraints in the form of monetary penalties. That is new constraints or new classes of constraint do not affect the space of possible locations of the intersection point, but rather are applied when deciding whether the new location is to be accepted. Reduced to the most basic situation, the permissible space for the new location of an intersection point is controlled solely by the geometric constraints.

Speed

For a given level of accuracy one of the major benefits of the simplicity of the location selection routines is speed. With stochastic optimisation it takes only a few minutes to complete a single optimisation run of moderate length, so that sets of twenty or so alternatives can be produced in a matter of hours. This operation does not require human involvement, so that the production of alignments under a wide variety of different conditions is relatively cheap, and leads to a number of other benefits.

Variable focus

Planning is essentially a cyclic process which gets more focused as it proceeds, and a gradually narrowing corridor is studied in increasing detail. To a large extent the power and flexibility of a model and its data requirements are inextricably linked. The greater the level of detail required from the model the more detail must be input, but data costs time and money. The speed and ease of running variations under stochastic optimisation means that a study may commence with relatively crude data and add more or finer detail incorporated later as the increased focus of the study justifies the cost of collecting the data.

Conflict resolution

The stochastic approach to optimisation with minimal absolute constraints imposed on the alignment also simplifies the identification and resolution of conflicts. In complex situations it is not uncommon for a particular formulation of the problem to be insoluble as a result of two or more constraints being mutually incompatible. In practical situations the information that a particular problem is insoluble is not in itself of much use, it merely poses the question of what can be done. Frequently many of the constraints imposed are desirable rather than necessary, and can be compromised to a greater or lesser extent to obtain an acceptable solution. If all constraints are absolute in their application it can be difficult to determine which are the critical constraints, let alone the severity of the problem and to what extent one or more may need to be relaxed. Employing cost penalties rather than absolute constraints simplifies the task of determining the nature, severity and the geographical location of the conflict, and can expedite the development of acceptable compromises.

Reliability

The time factor which was one of the principal reasons for investigating alignment optimisation in the first place has a second (indirect) effect which is not so widely recognised. Namely, it tends to discourage a thorough exploration of alternatives that are significantly from the original alignment, so that the cost of the final solution depends critically on the initial choice. That is, if the starting alignment is not close to the global optimum, the design process is likely to produce a final alignment whose costs cannot be reduced significantly by moving any one of the intersection points, but which nonetheless may be substantially more expensive than an alignment which runs up a different valley or passes a park on the other side. As a consequence, while the interactive approach to design is very good for cleaning up alignments and balancing subjective criteria against cost, it carries a significant risk of producing alignments that are only locally optimal.

Because a package based on stochastic optimisation can explore the whole transport corridor and can escape from local sub-optima, it will consistently produce low cost alignments that are uninfluenced by preconceptions of where the cheapest route should lie. If an alignment is cheap and satisfies the nominated constraints, the model has a high probability of finding it. If enough runs are undertaken all the better alignments should be represented.

CONCLUSIONS

Because it can explore the corridor thoroughly, stochastic procedures produce alignments that are as good as, if not better than, can be produced by other techniques. In all data sets employed to date the technique has been able to match or better alignments produced by conventional means. The savings in construction costs that can be achieved are hard to quantify because they depend to a large extent on the environment in the transport corridor. If the problem were to optimise an alignment across an empty plain no optimisation method could produce an improvement over conventional techniques. However, most problems are more complex, and stochastic optimisation can provide the greatest assistance in the most complicated cases. It is immaterial whether the complications arise from the complexity of the terrain or the nature and number of additional constraints and costs imposed on the route. Further, in many situations, it is not the location of the cheapest alignment that is important, but the provision of a selection of low cost alternatives that allow the planners to make an informed selection on the basis of qualitative factors.

The combination of fast optimisation with easy editing of data files makes it feasible to run optimisations with a range of geometric standards and environmental constraints, so that the costs of relaxing design standards or preserving specific environmental features can be realistically assessed. This approach can provide the planner with the material necessary for informed decisions, rather than try to quantify contentious costs within a program. For example, in the case of a private property or a nature reserve where intrusion could be expensive or controversial, the program would be run twice. The first time there would be no restrictions on the location of the alignment; the second time the property or reserve would be treated as an exclusion zone. The difference in construction costs between the two scenarios would then be balanced against the costs of acquisition, public relations etc to determine whether intrusion into the zone was a feasible option.

This approach has been incorporated in Align3D (Gipps 1992), a computer program developed by CSIRO Division of Building, Construction and Engineering, that has proved itself fast and flexible, able to generate low cost solutions in a very short timeAlign3D obtains its capabilities from a number of sources.

Because many aspects of a design are subjective, the acceptance of a particular design as 'good' depends very much upon the weight that a designer places on various criteria. Trying to meet these criteria within a computer package is misuse of resources. Instead, the development of program was based on leaving subjective issues in the domain of the human planners, and producing a range of alignments that met the objective criteria.

Optimising an alignment is complicated by the existence of many local sub-optima arising from the complexity of the land surface and the location of environmental features. Stochastic optimisation generates a range of altemative solutions, enables the construction of more realistic constraints, and provides the means by which an alignment can escape from local sub-optima.

Because the optimisation algorithm is independent of the cost function it is possible to handle complicated cost functions involving detailed construction requirements, and to amend these rapidly if necessary.

One factor which affected early work was the inadequacy of the computer hardware available and its effect on the choice of algorithms. Calculating construction costs for many variations of a route requires a lot of data which must be held in memory for fast computations. The lack of memory and lower speeds conditioned the choice of algorithms and limited the capabilities of the resulting packages.

REFERENCES

Chew, E.P., Goh, C.J. and Fwa, T.F. (1989) Simultaneous optimisation of horizontal and vertical alignments for highways, *Transportation Research* 23B, 315-329.

Gipps, P.G. (1992). ALIGN_3D: Package to optimise route alignment, *Road and Transport Research* 1, 50-59.

Goh, C.J. and Chew, E.P. (1988) Discrete and continuous models for computation of optimal vertical highway alignment, *Transportation Research* 22B, 399-408.

Nicholson, A.J., Elms, D.G. and Williman, A. (1976) A variational approach to optimal route location, *Highway Engineer* 23, 22-25.

Parker, N.A. (1977) Rural highway route corridor selection, *Transportation Planning Technology* 3, 247-256.

O'Brien, W.T. and Bennett, D.W. (1969) A dynamic programming approach to optimal route location, *Symposium on Cost Models and Optimisation in Road Location and Design*, London.

Trietsch, D. (1987) A family of methods for preliminary highway alignment, *Transportation Science* 21, 17-25.

Turner, A.K. (1978) A decade of experience in computer aided route selection, *Photogrammetry Engineering Remote Sensing* 44, 1561-1576.