# Decline design in underground mines using constrained path optimisation 

M. Brazil P.A. Grossman D.H. Lee J.H. Rubinstein D.A. Thomas<br>N.C. Wormald *


#### Abstract

In this paper we focus on the problem of optimising the design of an underground mine decline, so as to minimise the costs associated with infrastructure development and haulage over the life of the mine. A key design consideration is that the decline must be navigable by trucks and mining equipment, hence must satisfy both gradient and turning circle constraints. The decline is modelled as a mathematical network that captures the operational constraints and costs of a real mine, and is optimised using geometric techniques for constrained path optimisation. A deep understanding of the geometric properties of gradient and turning circle constrained paths has led to a very efficient procedure for designing optimal declines. This procedure has been automated in a new version of a software tool, Decline Optimisation Tool, DOT ${ }^{\text {TM }}$. A case study is described indicating the substantial improvements of the new version of DOT $^{\mathrm{TM}}$ over the earlier one.


Keywords: optimisation, underground mining, decline design

## 1 Introduction

In open-pit mine design, there are a number of widely used commercial optimisation software systems based on the Lerchs-Grossmann algorithm [1] and its improvements. In comparison, the usual method of deciding on a design for an underground mining network is based on a mining engineer's expertise and experience. This generally involves detailing a small number of feasible designs and choosing the "most suitable" one. Recently several researchers including Alford et al. [2] and Smith et al. [3] have shown that some underground mine design tasks such as stope layout and scheduling are amenable to optimisation techniques. Our group, based at The University of Melbourne, has been focusing on the opti-

[^0]misation of underground infrastructure and is developing software that will ultimately interface with the principal commercial modelling and visualisation packages used by the mining industry. The decline optimisation technique we describe is appropriate for a wide range of hard rock mines such as gold, silver, lead, zinc, copper or multimetal deposits where the orebodies are discrete and which are mined by stoping or caving methods.

The access infrastructure of an underground mine can be viewed as a network of interconnected declines and shafts that provides access to designated orebodies and a means of transporting ore from these zones to the mill. Finding an efficient layout for such a network is a difficult design problem. The mining engineer's initial solutions have the potential to be dramatically improved by systematic optimisation, due to the difficulty of designing efficient structures in three dimensions, the complexity of the possible connection schemes and the combinations of choices for operational constraints.

A decline is essentially a system of ramps and crosscuts (horizontal drives) that connects the access points (points which must be accessed for drilling and blasting operations) and draw points (from which the ore is drawn) to the surface portal or to a breakout from existing mine infrastructure. We shall assume that the locations and geometries of the stopes are given and that the stoping data has been used to determine the coordinates of the access points and draw points at each of the levels, as well as the tonnages of ore to be transported to the surface from each of the draw points.

There are a number of important physical constraints on the design of a decline. The decline must stand off from the orebody by some specified minimum distance to avoid possible sterilisation of the ore and to allow a minimum working length in the crosscuts. Another key design consideration is that the decline must be navigable by trucks and mining equipment; this limits the gradient and curvature of the decline. The aim in this paper is to minimise the cost of the decline, where the cost is a combination of both development and haulage costs, subject to these constraints.

In an earlier paper [4] the first version of the Decline Opti-
misation Tool, $\mathrm{DOT}^{\mathrm{TM}}$ (which we will refer to as DOT1), was described. In this paper, we describe the mathematical optimisation theory used to substantially improve the software tool. The heuristic methods used in DOT1 are replaced in the new version of the software tool, DOT2, by a method based on an understanding of exact solutions to a constrained 3 -dimensional path problem. This problem is described in the second section of the paper where we model the problem of designing an underground mine decline of minimum cost as a constrained path optimisation problem. In the third section, we develop a new theory of paths in 3-dimensional space that are optimal with respect to gradient and curvature constraints. These optimal paths which are used in DOT2 are shorter than the approximations obtained in DOT1. A case study is described in Section 6, indicating the significant improvements of the new version of $\mathrm{DOT}^{\mathrm{TM}}$ over the earlier one in terms of both the accuracy of the solution and the speed.

In designing the decline, typically there are no-go regions that must be avoided such as old mine workings. There must also be a standoff region for the decline from an orebody. These problems can be dealt with by adding barriers which the decline must not penetrate. In this paper we will limit the problem to designing a decline without barriers and we will assume homogeneous ground conditions. These preliminary strategic designs allow the decline a "free" path from the breakouts to stope access. In the earlier version, DOT1, a heuristic method of avoiding no-go regions had been implemented. The new version of the software tool, DOT2, that is described in this paper, does not yet have this capability but we plan to add barrier constraints in the future for comparison with the best case scenario. This comparison will be useful for mine management as it indicates the cost of certain decisions. For example it may influence a decision to add in a barrier or a decision on the distance for a standoff region.

## 2 The Mathematical Model of a Decline

We model the decline as a mathematical network that captures the operational constraints and costs of a mine. In the model the nodes of the network correspond to the orebody level access points and draw points and the surface portal (or breakout point from existing infrastructure) of the mine. The links in the network model represent the centrelines of ramps and drives.

The network model must include the navigability requirements imposed by the trucks and equipment to be used in the mine. The absolute value of the gradient of each ramp is constrained to be within a safe climbing limit for trucks, typically in the range $1: 9$ to $1: 6.5$. Hence the decline network is gradient constrained with a given maximum absolute value for the slope. In addition, there is a
minimum turning radius for curved ramps which is typically in the range $15-40 \mathrm{~m}$. These navigability constraints are significant factors in the optimal solution and to accommodate these the decline network is modelled as a gradient-constrained and curvature-constrained network.

A decline can be modelled as a network with a path topology if we ignore ventilation infrastructure and alternative means of egress. In particular, this applies to the case where there is a single orebody for a proposed new underground mine or an extension to an existing mine. We will focus on the fundamental problem of finding a least cost, navigable decline with a given path topology. In future work we plan to optimise mine layouts where there are a number of interconnected navigable declines in a tree structure.

Access to the orebody from the decline is via horizontal drives known as crosscuts. These connect the decline to the given access or draw points, which lie on a sequence of levels. Each crosscut should meet the decline at an angle of approximately 90 degrees for geomechanical stability. This is achieved by having a user-specified range of directions for the decline at each node. At each access level a set of candidate nodes, representing a discrete choice of junctions at which the crosscut can meet the decline, is specified. We refer to this set of nodes as a group. Each of these nodes has an associated fixed cost that is proportional to the length of crosscut and is dependent on the tonnage of ore to be hauled along the crosscut. There is a requirement that the decline goes through one node from each group. This provision of choice for the node locations provides design flexibility and optimisation opportunities.

We now model the optimal decline design problem as that of finding a smooth path in 3-dimensional space of minimum cost satisfying the following conditions:

1. It passes through one node from each of the groups of specified nodes at each level, in a given order; furthermore, for each node the path passes through, the bearing of the path at that node falls within a user-specified range of bearings;
2. At each point it has gradient at most $m$, where $m$ is a given constant;
3. At each point it has radius of curvature at least $r$ when projected into the horizontal plane, where $r$ is a given constant.

The cost of each link in the path is the sum of the development cost and haulage cost, where the development cost is proportional to length and the haulage cost is proportional to the length times the total tonnage hauled through that link (over the life of the mine or targeted
orebody). In particular, the cost of a link can be minimised by minimising its length, since the tonnage of ore to be hauled for each link is fixed. For a given path the total cost is the sum of the costs of all links in the path plus the sum of the costs associated with the selected nodes (one from each group).

The problem is solved by discretising the set of possible directions of the path at each of the nodes from the groups. This allows us to employ a bottom-up dynamic programming strategy. Suppose there are (up to) $n$ nodes in each group. The direction vectors at a node will have a limited number of directions entering (or leaving) the incident link depending on how close the node is to the orebodies. If the maximum number of directions at any node is $l$, then at each level or stage, we keep track of the (up to) $n l$ minimum cost subpaths constructed so far for the nodes at that level. Assuming we can construct a minimum cost link with given start and end-point directions in constant time, the path will be constructed in a time that is approximately quadratic in both the number of nodes at each level and directions at each node, and linear in the number of levels. More specifically, the algorithm runs in $O\left(k n^{2} l^{2}\right)$ time, where $k$ is the number of levels.

The dynamic programming framework essentially reduces the problem to one of efficiently minimising the length of a single link with given position and direction for each of the start and end points.

## 3 Link optimisation: Dubins paths and their extensions

An abstract solution to the problem of finding minimal paths in 3-dimensional space with given start and finish directions and a given minimum turning circle (but no gradient constraint) has been described in [5]. Unfortunately, this solution, as well as violating the gradient constraint, has the undesirable property of allowing the paths to have a continually varying gradient, whereas it is an industry requirement that the gradient on each link is both bounded and unchanging.

We approach the problem of minimising the cost of each link by considering the projected problem in the horizontal plane. Note that a path in the plane can be lifted to a path with a uniform gradient in 3-dimensional space. The length of this transformed 3-dimensional path depends only on the length of the planar path and the gradient. The transformed path will satisfy the gradient constraint if and only if the length of the path in the plane reaches a certain lower bound dependent on the vertical displacement between the end points of the link: if the vertical displacement is $z$ then the length is at least $z m^{-1}$, where $m$ is the given maximum gradient. Henceforth we use $B$ to denote this lower bound $z m^{-1}$, where $z$ is calculated with respect to whichever two end points are under
consideration.
Let $P$ be a path between two given directed points $p$ and $q$ in $\mathbf{R}^{2}$, the Cartesian plane. We call $P$ admissible if:

1. $P$ has a continuous first derivative and a piecewise continuous second derivative;
2. The tangents to $P$ at its start and end points coincide with the directions of $p$ and $q$ respectively;
3. The absolute curvature of $P$ is bounded above by a specified positive constant (which we will take to be 1 by choosing a suitable scaling).

The scaling condition in Item 3, above, implies that if the section of an admissible path is an arc then that section must have radius at least 1 .

The problem of finding the minimum cost link between two points in $\mathbf{R}^{3}$, 3-dimensional space, with given approach and departure directions and gradient and turning circle constraints can be solved by finding the shortest path in the plane between the projections of these points with these same direction and turning circle constraints and with a lower bound on the length, and then transforming this path back to $\mathbf{R}^{3}$. Specifically, we seek a minimum length admissible path, $P$, between two directed points $p$ and $q$ in the plane, such that the length of $P$ is at least $B$.

Dubins solved this planar problem without a lower bound on length in [6]. Our method for the planar problem where the lower bound is included is based on an extension of Dubins' work.

The main result of Dubins is as follows.

Theorem 3.1 [6] Given any two directed points, $p$ and $q$, in the plane, there exists an admissible path of minimum length from $p$ to $q$. Further, any such path must take one of the following forms:

- An arc with radius 1 and length less than $2 \pi$, followed by a line segment, followed by an arc with radius 1 and length less than $2 \pi$;
- A sequence of three arcs with radius 1 and with alternating senses (i.e., left-right-left or right-left-right), where the length of the middle arc is greater than $\pi$, and the length of each arc is less than $2 \pi$, thus less than a full circumference of a circle with radius 1.

Note that one or more of the arcs or line segments may be degenerate, in the sense that its length is zero. We refer to paths of the form given in Theorem 3.1 as Dubins paths. For any given pair of directed points, there may
be up to six Dubins paths. Using L, S and R to denote respectively a left turning arc, a (straight) line segment, and a right turning arc, we can identify each Dubins path by a unique descriptor, called its type: LSL, LSR, RSL, RSR, LRL and RLR. Three of the Dubins path types are illustrated in Fig. 1.


Figure 1: Three of the six types of Dubins paths; examples of the other three can be obtained from these by reflection through a horizontal axis.

Given directed points $p$ and $q$, a minimum length admissible path from $p$ to $q$ can be found simply by calculating the lengths of each of the Dubins paths from $p$ to $q$ and selecting the shortest path. If the shortest Dubins path has length at least $B$, then it is the solution to the original problem. Suppose now that the length of the shortest Dubins path is less than $B$. The approach we take here is to try to obtain an admissible path with length $B$ by continuously extending the shortest Dubins path. We will see that this is not always possible, in which case a solution, possibly with length greater than $B$, can be obtained using one of the other Dubins paths.

For the remainder of this paper, we use the term 'extension' of a path to mean a path obtainable by continuously deforming the original path, such that at all times during the deformation the path remains admissible. We now define two specific types of path extensions.

If a Dubins path, $P$, contains an arc with length at least $\pi$, then $P$ can be extended to an admissible path $P^{\prime}$ of any greater length in the manner illustrated in Fig. 2. In particular, any Dubins path of type LRL or RLR can be extended in this way, since the length of the middle arc of any such path is always greater than $\pi$. We refer to this type of extension as a parallel extension.

If the lengths of the arcs of a Dubins path are all less than $\pi$, then the situation is more complicated. Consider a Dubins LSL or LSR path, $P$, from $p$ to $q$. Then $P$


Figure 2: A parallel extension.



Figure 3: The two types of rolling extensions.
can be extended as follows. Let $C_{L}(p)$ and $C_{R}(p)$ denote the circles with unit radius that are tangent to the directed point $p$, on the left and right sides of $p$ respectively. (Thus, the first arc of $P$ is an $\operatorname{arc}$ of $C_{L}(p)$.) Similarly, let $C_{L}(q)$ and $C_{R}(q)$ denote the circles with unit radius that are tangent to the directed point $q$, on the left and right sides of $q$ respectively. It is helpful to imagine $P$ as an elastic band fixed at $p$ and $q$, and the four tangent circles as barriers that restrict the region in which $P$ can lie. Then $P$ can be extended either by "rolling" $C_{L}(p)$ clockwise around $C_{R}(p)$ to a new position $C_{L}(p)^{\prime}$, or by "rolling" $C_{R}(p)$ anticlockwise around $C_{L}(p)$, keeping the other three circles fixed, as shown in Fig. 3. We refer to these two types of extensions as rolling extensions.

In the configurations depicted in Fig. 3, $P$ could be extended indefinitely by rolling one circle sufficiently far around the other. However, as far as the present application is concerned, it makes more sense to extend $P$ in this manner only until one of the arcs achieves a length of $\pi$, at which stage a parallel extension can be applied if necessary. We call a path infinitely extendible if it has arbitrarily long extensions.

For some configurations of $p$ and $q$, a rolling extension can be carried out only until the extended path $P^{\prime}$ reaches a local maximum admissible path. By a local maximum admissible path we mean an admissible path that is a local maximum under continuous deformations of the path (that always remain admissible); in other words, a path such that there does not exist an arbitrarily small deformation to a longer path within the space of admissible paths. If $P^{\prime}$ achieves the required length $B$ before reaching a local maximum admissible path, then the local maximum causes no difficulties. However, if the length of the local maximum path is less than $B$, then this strat-


Figure 4: An example of a path that is not infinitely extendible.
egy does not work. Applying the other type of rolling extension may prove successful, but this is not always the case. Fig. 4 depicts a situation where the path $P^{\prime}$ "gets stuck" between the four tangent circles. For example, the lower path in broken lines in the figure can be obtained by rolling $C_{L}(p)$ around the circumference of $C_{R}(p)$ until it meets $C_{R}(q)$. If we continue to roll $C_{L}(p)$ further around the circumference of $C_{R}(p)$, then under a continuous deformation the path from $p$ is forced to enter the circle $C_{R}(q)$, and cannot meet up with $q$ at the correct bearing without either breaking the turning circle constraint or substantially increasing the length of the path (by first leaving the circle $C_{R}(q)$, and then looping around to meet $q$ at the correct angle). It follows that the two paths shown as broken lines are both local maxima and cannot be lengthened by small deformations under the admissibility conditions. The paths are local maxima but not global maxima (where global maximum simply means the longest admissible path with no condition on continuous deformations).

It is important to characterise local maxima. We first note that if $P$ is a local maximum admissible path in the plane between two directed points then $P$ contains no twice differentiable points with absolute curvature less than 1. This follows from the observation that the length of a segment of the path with absolute curvature less than 1 can be increased (while remaining admissible) by the perturbation of a unit circle tangent to an interior point of such a segment, resulting in replacing part of the segment by three unit circle arcs.

It follows that $P$ must be a sequence of unit circle arcs with alternating senses, which we denote by a sequence of C's. So, for example, CCC represents a path of type LRL or RLR. The key theorem is as follows.

Theorem 3.2 Given two directed points, $p$ and $q$, in the plane, let $P$ be an admissible path from $p$ to $q$. If $P$ is a local maximum then $P$ is of the form $C C C$ (or a degeneracy) where the angle around each circle is less than $\pi$ and the sum of the two angles around the outer two circles minus that around the inner circle is less than $\pi$.

The proof of the theorem, which is fairly straightforward, involves showing that a path of the form CCCC is not a local maximum. Details of the proof are planned to appear in a future paper.


Figure 5: A flow chart showing how to construct a minimum cost admissible path from $p$ to $q$.

## 4 A Minimal Path algorithm

We now outline an algorithm for constructing a single minimum cost link in 3-dimensional space with given approach and departure directions and gradient and turning circle constraints. The algorithm is presented as a flow chart in Figure 5. Theorem 3.1 allows one to calculate all the possible Dubins paths and hence to identify the shortest, $P_{1}$, and second shortest, $P_{2}$, in Step 1. Theorem 3.2 characterises the local maximum admissible paths that can be obtained from $P_{1}$, allowing one to calculate $L_{\text {max }}$ in Step 3. The output of the algorithm indicates how to construct the minimum cost link. The final decline is built from these minimum cost links via dynamic programming, as discussed in Section 2, and is minimum up to the degree of discretisation of possible positions and angles at the nodes.

The algorithm first solves the planar path problem with lower bound $B$ considered in Section 3. For a given link, let $p$ and $q$ be the projections in the horizontal plane of its
two endpoints. Apart from certain degenerate configurations which we describe below, at most one of the Dubins paths from $p$ to $q$ can fail to be infinitely extendible.

If two or more Dubins paths coincide because of the presence of degenerate arcs or line segments, then the shortest and second shortest Dubins paths may be the same. If this occurs, then $P_{2}$ is the second shortest distinct Dubins path in Step 5.

The algorithm fails if the minimum Dubins path from $p$ to $q$ takes one of the following degenerate forms: a line segment with length less than $4 a$ where $a$ is the unique number satisfying the equation $\cos a=a$, or an arc with length less than $\pi / 2$. In these cases, the algorithm must be modified by taking $P_{2}$ to be the path obtained by appending a circle, i.e., an arc of length $2 \pi$, to $P_{1}$. Note that the algorithm assumes that if $P_{1}$ is not infinitely extendible then $P_{2}$ is infinitely extendible. A proof of this is planned to appear in a future paper.

Observe that none of the 3 -dimensional paths need to be constructed in the course of running the algorithm; rather, the algorithm returns a Dubins path type together with information on whether and how the Dubins path is to be extended. Once the length $L_{p q}$ of the optimal planar path between $p$ and $q$ has been computed, then it is easily seen that the length of the corresponding path in 3dimensional space (furnished with a constant gradient) is $\sqrt{L_{p q}^{2}+\left(z_{p}-z_{q}\right)^{2}}$, where $z_{p}$ and $z_{q}$ are the $z$ coordinates of the two endpoints of the link. This strategy of computing lengths and recording the type of Dubins path, rather than constructing each 3-dimensional path during the dynamic programming, ensures that ultimately only the links that are actually needed for the optimal decline path are constructed.

Finally, we comment briefly on the problem of constructing the links in the optimal decline path. A link is represented by a list of line segments and arcs, where each line segment is parameterized by its start and end points, and each arc is parameterized by its centre, start angle and turn angle. Constructing a Dubins path and a parallel extension are straightforward, but constructing a rolling extension is more difficult, and generally requires the use of an iterative procedure.

Once a planar path has been constructed for a given link it is converted to a path in 3-dimensional space by giving it the correct constant gradient: $\pm\left(z_{p}-z_{q}\right)\left(L_{p q}\right)^{-1}$.

## 5 Further Practical Constraints on the Path Optimisation Problem

Both DOT1 and DOT2 design the decline link by link. This may result in some links having gradients less than the maximum gradient. For two given adjacent access points, it may not be possible to construct a link of max-
imum gradient simply because of their physical distance apart. If a decline with a constant gradient throughout is required, and there is at least one link with less than maximum gradient in the optimal solution, then in the new decline every link would have to have less than maximum gradient. A decline with constant gradient can be achieved by rerunning DOT2 using successively smaller maximum gradients until a suitable decline is generated. This new decline, however, may be substantially longer than the optimal solution under the original maximum gradient.

In discussion with industry, it has become clear that there should be an aim to reduce the occurrence of adjacent arcs with opposite senses (mathematically, points of inflexion) such as occurs in an LRL or an RLR path; such features are usually avoided by mining engineers, as far as possible, because of the physical problem of reversing the direction of camber of the road surface at the inflexion point, as well as the difficulty it causes for the drivers of the ore trucks. We have modified the path extension methods to take account of the opposing arcs constraint, by ensuring that there is a straight of length at least 10 metres between such arcs, and incorporated the modifications into $\mathrm{DOT}^{\mathrm{TM}}$. In a later paper, the question of the optimality of the algorithm with the opposing arcs constraint will be examined carefully. This opposing arcs constraint is included in the DOT2 software tool demonstrated in the case study described below.

## 6 Case Study

DOT ${ }^{\mathrm{TM}}$ has been tested mainly in design tasks for various Australian and New Zealand mines operated by Newmont Australia Limited, our collaborative research partner. This has been particularly valuable in refining the features in $\mathrm{DOT}^{\mathrm{TM}}$ to match both operational and strategic design needs; the new algorithm, DOT2 incorporates many refinements over the heuristic DOT1 - mainly the improved algorithm outlined in Section 4. This algorithm uses optimal paths which give significant savings in cost.

In June 2006 we were offered the chance to compare our design with one developed by an experienced mine consultant. The design was required to span 18 given access or draw points with a decline maximum gradient of 1:7 and a minimum turning radius of 25 metres. We were able to compare our DOT1 and DOT2 designs against the engineer's design. Fig. 6 compares the designs in a composite representation. The heuristic algorithm can be set at different levels dependent on the level of accuracy. It took about 20 minutes at a reasonable level of accuracy to find a decline of 1883 metres in length. The new algorithm is optimising for the given input and took only a few seconds to find a design of length 1771 metres. The engineer's original design was 1964 metres in length. Thus the new algorithm saved about $10 \%$ over the orig-


Figure 6: A comparison of the engineer's original design with the declines generated by DOT1 and DOT2.
inal design and significantly outperformed the heuristic in time and total decline length. Since a metre of decline development currently costs about AU $\$ 4,000$, the development savings alone from DOT2 are of the order of AU $\$ 772,000$, compared to the original engineer's design. Over the life of a mine the corresponding savings in haulage, ventilation and other operational costs may double this, leading to overall savings of the order of $\mathrm{AU} \$ 1.5$ million.

## 7 Conclusions

We have presented a new algorithm for designing underground mine declines so as to optimise the associated life-of-mine costs, based on a mathematical analysis of the properties of minimum length curvature-constrained paths. This approach, implemented in DOT2, substantially improves on previous methods, in terms of both speed and accuracy. The efficiency of DOT2 allows alternative decline designs to be generated and displayed in seconds. The automation of this component of the design process allows the mining engineer to consider and explore alternative development scenarios.

## 8 Acknowledgments

This research has been conducted with financial support from the Australian Research Council and Newmont Asia Pacific Limited via a Linkage Collaborative

Research Grant. Nicholas Wormald is also supported by the Canada Research Chairs program. We acknowledge the contributions of Andrew Fox, Nadine Wetzel and Simon Kusabs of Newmont Asia Pacific Limited for their work in clarifying Newmont's software requirements, providing valuable feedback on prototype versions of DOT2, and supplying us with real mine data to test DOT2.

## References

[1] H. Lerchs and I.F. Grossmann: "Optimum Design of Open-pit Mines", Trans. Canad. Inst. Min. Metall., 1965, 68, 17-24.
[2] C. Alford, M. Brazil and D.H. Lee: "Optimisation in underground mining", in J.P. Miranda, A. Weintraub, C. Romero, T. Bjørndal, and R. Epstein (eds.), Handbook of Operations Research in Natural Resources, 2007, 561-578 (Springer: Berlin).
[3] M.L. Smith, I. Sheppard and G. Karunatillake: "Using MIP for strategic life-of-mine planning of the lead/zinc stream at Mount Isa Mines", APCOM 2003 Conference, Cape Town, South Africa, May 2003, 1-10.
[4] M. Brazil, D.H. Lee, M. Van Leuven, J.H. Rubinstein, D.A. Thomas, and N.C. Wormald: "Optimising Declines in Underground Mines", Min. Technol.
(Trans. Inst. Min. Metall. A), 2003, 112, A164A170.
[5] H.J. Sussman: "Shortest 3-dimensional paths with prescribed curvature bound", Proc. 34th Conf. on Decision and Control, IEEE, New Orleans, USA, December 1995, 3306-3312.
[6] L.E. Dubins: "On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents", Amer. J. Math., 1957, 79, 497-516.


[^0]:    *Addresses: M. Brazil, P.A. Grossman, D.A. Thomas: ARC Special Research Centre for Ultra-Broadband Information Networks (CUBIN), an affiliated program of National ICT Australia, Department of Electrical and Electronic Engineering, The University of Melbourne, Victoria 3010 Australia; D.H. Lee, J.H. Rubinstein: Department of Mathematics and Statistics, The University of Melbourne, Victoria 3010 Australia; N.C. Wormald: Department of Combinatorics and Optimization, University of Waterloo, Waterloo ON N2L 3G1, Canada. Contact Email: D.A. Thomas: d.thomas@ee.unimelb.edu.au

