Journal of Artificial Intelligence Research 59 (2017) 543-563 Submitted 01/17; published 08/17

Research Note

The Length of Shortest Vertex Paths in Binary Occupancy Grids Compared to Shortest r-Constrained Ones

Patrick Chisan Hew PATRICK.HEW@DEFENCE.GOV.AU
Defence Science and Technology Group
Department of Defence (Australia)
HMAS Stirling, Garden Island WA 6168, Australia

Abstract

We study the problem of finding a short path from a start to a goal within a two-dimensional continuous and isotropic terrain that has been discretized into an array of accessible and blocked cells. A classic approach obtains a *grid path* where each step is along the edge of an accessible cell or diagonally across one. Grid paths suffer from 'digitization bias' – even if two locations have line-of-sight, the minimum travelling cost between them can be greater than the distance along the line-of-sight. In a *vertex path*, steps are allowed from a cell corner to any other cell corner if they have line-of-sight. While the 'digitization bias' is smaller, shortest vertex paths are impractical to find by brute force. Recent research has thus turned to methods for finding short (but not necessarily shortest) vertex paths. To establish the methods' potential utility, we calculate upper bounds on the difference in length between the shortest vertex paths versus the shortest r-constrained ones where an r*-constrained path* consists of line segments that each traverse at most r rows and at most r columns of cells. The difference in length reduces as r increases – indeed the shortest vertex paths are at most 1 percent shorter than the shortest 4-constrained ones. This article will be useful to developers and users of short(est) vertex paths algorithms who want to trade path length for improved runtimes in a predictable manner.

1. Introduction

We study the problem of finding a short path that a vehicle can traverse from a start to a goal within a two-dimensional continuous terrain where locations are either habitable or not. An example is navigating a watercraft from port to port while staying in a sufficient depth of water. Similar problems arise in mobile robotics and in the realistic-looking routing of entities in computer games. We assume that the terrain has been discretized into a two-dimensional array of accessible and blocked cells (a *binary occupancy grid* – Figure 1.a into Figure 1.b). The cells' corners are declared to be the vehicle's *feasible locations.* We assume that if a line-of-sight between feasible locations passes only through accessible cells then the vehicle can traverse that line-of-sight in the original terrain. We further assume that when travelling in accessible terrain, the vehicle can move at the same speed in any direction (motion is isotropic). Finally, we assume that the vehicle's dimensions and turning circle are small compared to the cells so it can 'squeeze through the diagonal' between cells that touch at corners but not faces. The problem's formulation aligns with readily-available data sets; for example the ETOPO data set (Amante & Eakins, 2009) provides the altitude above sea level of the Earth's surface using cells that have side lengths of 1 arc-minute.

Unfortunately, while it is easy to represent terrain as a two-dimensional array of accessible and blocked cells, the search within this representation for a shortest path from start to goal is not trivial. The *true shortest paths* between any two locations are the ones that minimize the travelling cost between them (this holds for any locations in the terrain, not just the feasible ones). In isotropic terrain, travelling cost is directly proportional to path length. If the locations have line-of-sight then the true shortest path between them is along that line-of-sight. So to approximate a true shortest path from a start to a goal it is sufficient to find a feasible location near to the start, a feasible location near the goal, and obtain a *shortest vertex path* between those two locations. Here, a *vertex path* is a path through the visibility graph of the feasible locations; by definition, nodes are adjacent in the *visibility graph* if they have line-of-sight to each other (Figure 1.c). Thus in principle, to obtain a shortest vertex path we need only search the visibility graph of feasible locations. But in practice, the graph is expensive to compute, and too large to search by brute force.

Consequently it is common to consider the paths through the *grid graph* on the feasible locations, declaring nodes as adjacent if they share an accessible cell (same edge or diagonally across). Hence in a *grid path*, each step is along the edge of an accessible cell or diagonally across one; this is equivalent to 8-connected motion (Figure 1.d). Grid paths suffer from 'digitization bias' (Tsitsiklis, 1995): even if two nodes have line-of-sight, the minimum travelling cost between them can still be greater than the distance along the line-of-sight.

Algorithms for planning short (but not necessarily shortest) vertex paths have therefore received much attention (see Uras & Koenig, 2015, for a review of such planners including so-called 'any angle' algorithms). As there is typically a tradeoff between path length and runtime, developers and users of short vertex path planners may need to tune their algorithms to obtain 'good enough' path lengths within an acceptable runtime. We inform such efforts by considering paths that are constrained as follows:

> **Definition**: A path is an *r-constrained path* if it consists of line segments that each traverse at most r rows and at most r columns of cells.

We will then prove the following:

> **Main result**: Let $\mathcal{P}$ be a vertex path from a start to a goal and r be any positive integer. Then there exists an r-constrained path $\mathcal{P}'$ from the same start to the same goal such that
>
> $$\frac{\|\mathcal{P}'\|_2 - \|\mathcal{P}\|_2}{\|\mathcal{P}'\|_2} \leq 1 - \cos\left(\frac{\operatorname{arccot}(r)}{2}\right)$$
>
> where $\|\cdot\|_2$ is the Euclidean length. The result holds in particular when $\mathcal{P}$ is a shortest path. Moreover the bound applies to the shortest r-constrained path from the start to the goal, or any other path $\mathcal{P}''$ that is of the same length or shorter than $\mathcal{P}'$.

That is, rather than finding a shortest vertex path, it may 'good enough' to find a shortest r-constrained one. Notably, the shortest vertex paths are at most 1 percent shorter than the shortest 4-constrained paths. Researchers can use the results of this paper to establish whether a given short vertex path planner is suitable for their application, to anticipate the improvements that could accrue and to pursue improvements in algorithms.

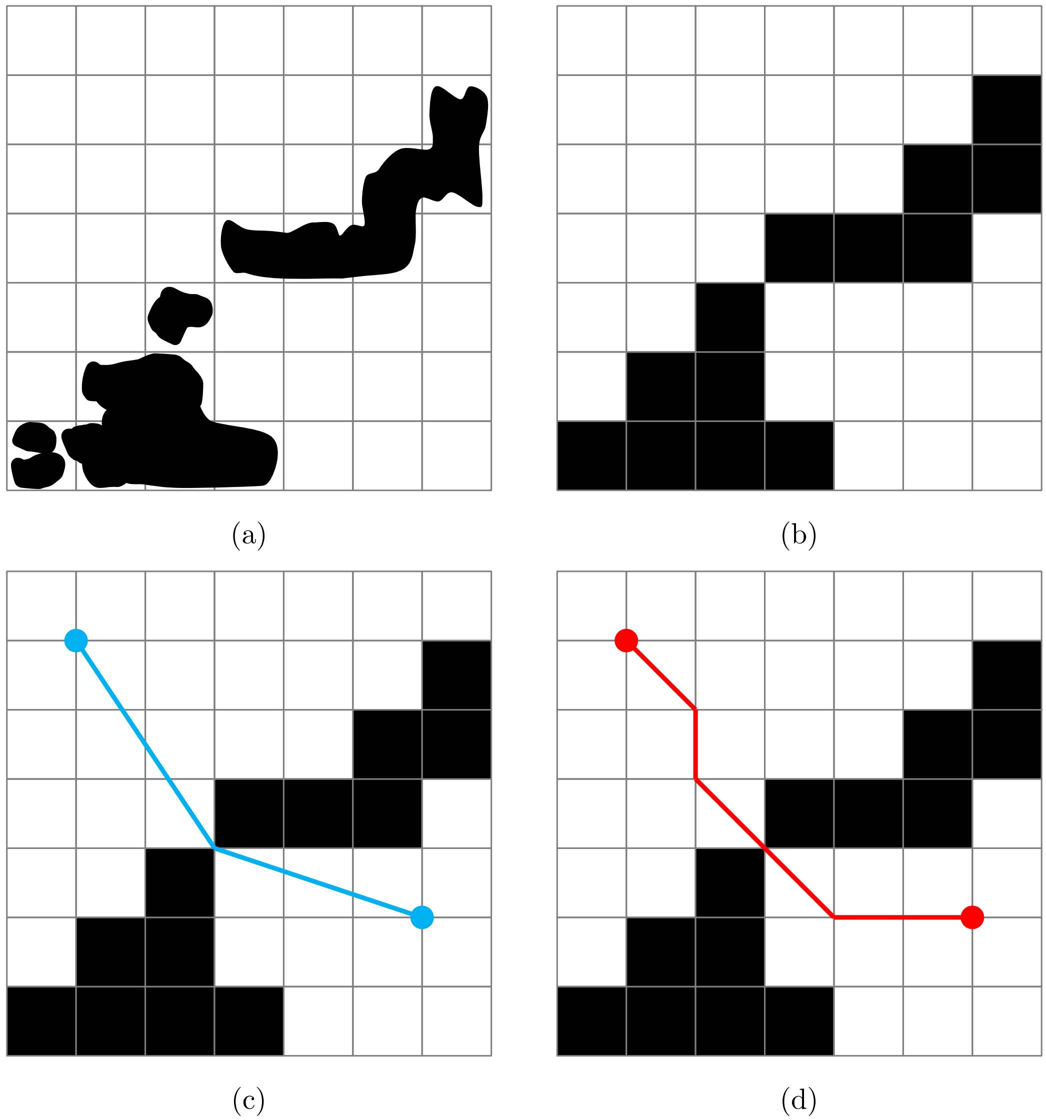

Figure 1: Representing the terrain to find short paths:
a) Locations are either habitable or not (white is accessible, black is forbidden).
b) Discretize into a two-dimensional array of accessible and blocked cells.
c) Vertex paths proceed along feasible locations that have line-of-sight.
d) Grid paths suffer from 'digitization bias'.

Previous investigations have focussed on the length of shortest grid paths (1-constrained paths) versus shortest vertex paths. Nash's comprehensive study (2012, also see the overview in Nash & Koenig, 2013) considered grid paths on regular grids in two and three dimensions. His work was subsequently expanded by Bailey et al. (2015) in the two-dimensional case to cover both 4- and 8-connected paths, with feasible locations at both cell corners and cell centers (they also studied the lengths of shortest vertex paths versus true shortest paths).

2. Reduction in Lengths of Shortest Paths

We will be working with lengths and gradients of line segments. For these concepts to be well-defined, we define a *grid-aligned coordinate system* as a coordinate system where one axis is parallel to the rows of the array, the other axis is parallel to the columns, the origin is at a cell corner, and the cells are declared to be of unit size. Note that for any line segment, we can choose a grid-aligned coordinate system such that the line segment has a gradient between zero and one inclusive (Specifically: Given a line segment between vertices on a binary occupancy grid, choose a pair of axes that puts the line segment in the positive quadrant with a non-negative gradient. Then if the gradient is greater than one, swap the axes). We will use this fact to simplify our proofs.

We consider finite *sequences* of feasible locations $\{s_k\}_k$. Sequences of two points $\{A, B\}$ may be abbreviated to AB. A *path* is a sequence in which s_k has line of sight to s_{k+1} for all k. We write $\partial(s, w)$ for the gradient of the line segment from s to w, $\|\cdot\|_2$ for the Euclidean norm (ℓ_2 distance) and $\|\cdot\|_\infty$ for the Chebyshev norm (ℓ_∞ distance). For sequence $\mathcal{Q} = \{s_k\}_k$ and for any norm $\|\cdot\|$ we declare $\|\mathcal{Q}\| = \sum_k \|s_{k+1} - s_k\|$. Given this notation, we may restate the definition of r-constrained sequences:

Definition (r-constrained sequences, restated)**.** Let r be a positive integer. A sequence $\mathcal{Q} = \{s_k\}_k$ is *r-constrained* if $\|s_{k+1} - s_k\|_\infty \leq r$ for all k. Equivalently, there are at most r rows and at most r columns of cells between s_k and s_{k+1} for all k.

The crux of our approach is that for any shortest vertex path and for any positive integer r, there exists an r-constrained path with the same start and goal. We can study the length of that r-constrained path. This will yield an upper bound on the differences in length between the shortest vertex paths and the shortest r-constrained paths. Hence we obtain an upper bound on the *inefficiency* which we define as follows:

Definition (Inefficiency)**.** Let $\mathcal{Q}$ be a sequence. Let $\mathcal{Q}'$ be a sequence that has the same start and goal as $\mathcal{Q}$ but has a longer Euclidean length. We call the relative (percentage) amount that $\mathcal{Q}$ is shorter than $\mathcal{Q}'$ the *inefficiency of $\mathcal{Q}'$ with respect to $\mathcal{Q}$*, calculated as

$$\bar{\eta}(\mathcal{Q}'; \mathcal{Q}) = \frac{\|\mathcal{Q}'\|_2 - \|\mathcal{Q}\|_2}{\|\mathcal{Q}'\|_2}$$

We will develop our results for line segments (paths of two points) and then apply them to paths as a straightforward extrapolation. We will need the following technical result that will be put into context in due course.

Lemma 1. Let $\mathcal{Q}$, $\mathcal{Q}'$, $\mathcal{Q}''$ be sequences with the same start and goal locations and suppose that $\|\mathcal{Q}\| \leq \|\mathcal{Q}'\| \leq \|\mathcal{Q}''\|$. Then

$$\frac{\|\mathcal{Q}'\| - \|\mathcal{Q}\|}{\|\mathcal{Q}'\|} \leq \frac{\|\mathcal{Q}''\| - \|\mathcal{Q}\|}{\|\mathcal{Q}''\|}$$

Proof.

$$\frac{\|\mathcal{Q}'\| - \|\mathcal{Q}\|}{\|\mathcal{Q}'\|} = 1 - \frac{\|\mathcal{Q}\|}{\|\mathcal{Q}'\|} \leq 1 - \frac{\|\mathcal{Q}\|}{\|\mathcal{Q}''\|} = \frac{\|\mathcal{Q}''\| - \|\mathcal{Q}\|}{\|\mathcal{Q}''\|}$$

□

2.1 An Upper Bound on Inefficiency from Differences in Gradient

Suppose that we have a sequence $\mathcal{Q}$ from A to B. What is the worst that the inefficiency can be with respect to AB? That is, can we obtain an *upper bound on the inefficiency*? We will see that an upper bound exists, based on the following quantity:

Definition (Maximum difference in angle)**.** Let $\mathcal{Q} = \{s_k\}_k$ be a sequence with increasing x coordinates (under some grid-aligned coordinate system). Then

$$\gamma(\mathcal{Q}) = \arctan\left(\max_k \partial(s_k, s_{k+1})\right) - \arctan\left(\min_k \partial(s_k, s_{k+1})\right)$$

In words: $\gamma(\mathcal{Q})$ is the difference in angle between the steepest and shallowest line segments of $\mathcal{Q}$. We measure γ in radians.

The idea is to consider $\mathcal{Q}$ as a sequence of line segments and rearrange them so that they are joined end-to-end by increasing gradient. Then $\gamma(\mathcal{Q})$ is the angle from the first line segment to the last one. Indeed by extending those line segments until they intersect, we obtain a triangle that has $\gamma(\mathcal{Q})$ as an exterior angle. The upper bound on inefficiency comes from analyzing that triangle.

We first establish that we can rearrange $\mathcal{Q}$ without losing anything important.

Lemma 2. Let $\mathcal{Q}$ be a sequence of locations with increasing x coordinates (in some grid-aligned coordinate system). Sort the line segments of $\mathcal{Q}$ by increasing gradient, then join them end-to-end. Then the resulting sequence $\mathcal{Q}'$ has the same start, end, and length as $\mathcal{Q}$.

Proof. By construction, $\mathcal{Q}'$ has the same start as $\mathcal{Q}$. The line segments' lengths are preserved so $\mathcal{Q}'$ has the same length as $\mathcal{Q}$. The gradients are also preserved so the cumulative displacement from the start of $\mathcal{Q}$ is also preserved. Thus $\mathcal{Q}'$ has the same end as $\mathcal{Q}$. □

Remark. $\mathcal{Q}$ may be a path but the resulting sequence $\mathcal{Q}'$ need *not* be a path. As we are interested in $\mathcal{Q}'$ purely for its shape, it does not matter if it passes through blocked cells.

Remark. This article generalizes the approach taken by Nash (2012), to handle line segments at a finite number of gradients (Nash, 2012, assumed two gradients only).

By extending the first and last line segments until they intersect, we form a triangle.

Lemma 3. Let $\mathcal{Q}$ be a sequence from A to B such that (under some grid-aligned coordinate system) the x coordinates of $\mathcal{Q}$ are increasing, the gradients of the sequence's line segments take on at least two distinct values, the gradients of successive line segments are increasing, and $\gamma(\mathcal{Q}) < \pi$. Extend the first and last line segments of $\mathcal{Q}$ into lines and construct C as the intersection of those lines (the conditions on $\gamma(\mathcal{Q})$ make C well-defined). Put $\mathcal{Q}'' = \{A, C, B\}$. Then $\|\mathcal{Q}\| \leq \|\mathcal{Q}''\|$.

Proof. If the line segments of $\mathcal{Q}$ have two gradients then write $\mathcal{Q} = \{A, s_1 \ldots s_{m-1}, C, s_{m+1} \ldots s_{n-1}, B\}$ where $s_0 = A$, $s_n = B$, $s_m = C$ for some positive integers m, n. Then $\|\{s_0 \ldots s_m\}\| = \|AC\|$ and $\|\{s_m \ldots s_n\}\| = \|CB\|$ hence $\|\mathcal{Q}\| = \|\mathcal{Q}''\|$. Otherwise write $\mathcal{Q} = \{A, s_0 \ldots s_{n-1}, s_n, B\}$ so that $\mathcal{Q}$ consists of $n+2$ line segments with strictly increasing gradients, for some positive integer n. Extend the last and third-last line segments into lines, let C' be their intersection (they are not parallel so C' is well-defined), and let $\mathcal{Q}' = \{A, s_0 \ldots s_{n-1}, C', B\}$. Then $\|\mathcal{Q}\| \leq \|\mathcal{Q}'\|$ by the triangle inequality on $s_{n-1}C's_n$. Repeating this process eventually yields $\mathcal{Q}' = \{A, C, B\}$ and $\|\mathcal{Q}\| \leq \|\mathcal{Q}''\|$, as required. □

Remark. The point C need *not* be a cell corner. We are only interested in triangle ABC for its dimensions so it does not matter if it is a sequence or not.

We now analyze the geometry of the triangle.

Lemma 4. For $c > 0$ and $0 \leq \gamma < \pi$ fixed, let $a, b > 0$ such that the triangle ABC of Figure 2 is well-formed. Put $\Delta = 1 - \cos\frac{\gamma}{2}$. Then $\dfrac{a+b-c}{a+b} \leq \Delta$ for all valid a, b, with equality when $a = b$.

Proof. We have $\beta = \gamma - \alpha$ and thus $0 \leq \alpha \leq \gamma$. By the sine rule

$$\frac{a}{\sin\alpha} = \frac{b}{\sin\beta} = \frac{c}{\sin(\pi - \gamma)}$$

so

$$\begin{aligned} \frac{a+b-c}{a+b} &= 1 - \frac{\frac{\sin\gamma}{\sin\beta}b}{\frac{\sin\alpha}{\sin\beta}b + b} \\ &= 1 - \frac{\sin\gamma}{\sin\alpha + \sin(\gamma - \alpha)} \end{aligned}$$

Construct $f(\alpha) = 1 - \dfrac{\sin\gamma}{\sin\alpha + \sin(\gamma - \alpha)}$. We maximize f as a function of α. Note that if triangle ABC is well-formed then $0 < \alpha < \gamma$. Now f is differentiable on its domain; indeed

$$\frac{df}{d\alpha} = \sin\gamma \frac{\cos\alpha - \cos(\gamma - \alpha)}{(\sin\alpha + \sin(\gamma - \alpha))^2} \begin{array}{ll} > 0 & \text{for } 0 \leq \alpha < \frac{\gamma}{2} \\ = 0 & \text{for } \alpha = \frac{\gamma}{2} \\ < 0 & \text{for } \frac{\gamma}{2} < \alpha \leq \gamma \end{array}$$

That is, f increases to its maximum at $\alpha = \frac{\gamma}{2}$ and decreases thereafter so

$$\frac{a+b-c}{a+b} \leq f(\tfrac{\gamma}{2}) = 1 - \frac{\sin\gamma}{2\sin\frac{\gamma}{2}} = 1 - \cos\tfrac{\gamma}{2} = \Delta$$

Moreover if $\alpha = \frac{\gamma}{2}$ then $\beta = \frac{\gamma}{2}$ and thus $a = b$. □

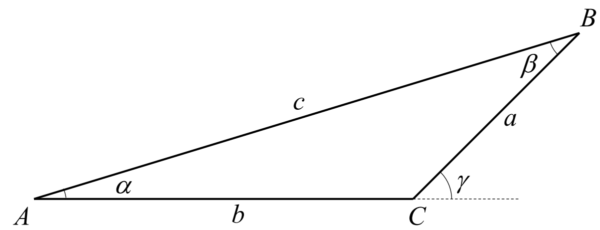

Figure 2: Proof of Lemma 4: For any fixed γ, the relative (percentage) difference between $a+b$ and c is maximized when $a=b$.

Putting everything together yields our upper bound on inefficiency.

Proposition 1 (Upper bound on inefficiency)**.** Choose a grid-aligned coordinate system such that the line segment AB is accorded a positive gradient (without loss of generality). Suppose that $\mathcal{Q} = \{s_k\}_k$ is a sequence from A to B that has increasing x coordinates and $\gamma(\mathcal{Q}) < \pi$. Then

$$\frac{\|\mathcal{Q}\|_2 - \|AB\|_2}{\|\mathcal{Q}\|_2} \leq 1 - \cos\left(\frac{\gamma(\mathcal{Q})}{2}\right)$$

Proof. Form $\mathcal{Q}$ into $\mathcal{Q}'$ by sorting its line segments by increasing gradient as per Lemma 2. Extend the first and last line segments of $\mathcal{Q}'$ into lines and construct C as the intersection of those lines (the conditions on $\gamma(\mathcal{Q})$ make C well-defined). Put $\mathcal{Q}'' = \{A, C, B\}$ and note that $\|\mathcal{Q}'\| \leq \|\mathcal{Q}''\|$ by Lemma 3. Then

$$\begin{aligned}
\frac{\|\mathcal{Q}\|_2 - \|AB\|_2}{\|\mathcal{Q}\|_2} &= \frac{\|\mathcal{Q}'\|_2 - \|AB\|_2}{\|\mathcal{Q}'\|_2} && \text{by Lemma 2} \\
&\leq \frac{\|AC\|_2 + \|CB\|_2 - \|AB\|_2}{\|AC\|_2 + \|CB\|_2} && \text{by Lemma 1} \\
&\leq 1 - \cos\left(\frac{\gamma(\mathcal{Q})}{2}\right) && \text{by Lemma 4}
\end{aligned}$$

□

2.2 Existence of r-constrained Paths with Bounded Differences in Gradient

In the previous section, we showed that if $\mathcal{P}$ is a path from A to B then there is an upper bound on the inefficiency of $\mathcal{P}$ with respect to AB that can be calculated from $\gamma(\mathcal{P})$ (the difference in angle between the steepest and shallowest line segments of $\mathcal{P}$). We now show that if A has line-of-sight to B then there exists an r-constrained path $\mathcal{P}$ from A to B *and moreover* there is an upper bound on $\gamma(\mathcal{P})$. The two results, put together, will yield the main result for this paper.

The centerpiece of this section is Algorithm 1: If A has line-of-sight to B then Algorithm 1 will construct an r-constrained path $\mathcal{P}$ from A to B and the line segments making up $\mathcal{P}$ will have gradients that will allow us to bound $\gamma(\mathcal{P})$. As notation, let u_x, u_y denote the x, y coordinates of a given location u. The following definitions also apply:

Definition (Flooring of a line segment)**.** Suppose that $A = (0,0)$ and $B = (q,p)$ where $0 < p < q$ (in some grid-aligned coordinate system). Let X be a subsequence from $0 \dots q$ (in increasing order). We define the *flooring of AB over X* as

$$\lfloor AB \rfloor_X = \left\{ \left(x, \lfloor \tfrac{p}{q} x \rfloor \right) \right\}_{x \in X}$$

For brevity we write $\lfloor AB \rfloor = \lfloor AB \rfloor_{0 \dots q}$.

Definition (Farey sequences and Farey pairs, Weisstein, 2017)**.**

1. We write $\mathcal{F}_n$ for the *Farey sequence of order n*, namely the irreducible rational numbers $\frac{a}{b}$ where $0 \le a \le b \le n$ (and a, b are coprime), arranged in increasing order.

2. $\frac{a}{b}$ and $\frac{c}{d}$ are a *Farey pair* if they are neighbours in a Farey sequence. Moreover if $\frac{a}{b} < \frac{c}{d}$ then $bc - ad = 1$. If $\frac{a}{b} < x < \frac{c}{d}$ for some x then we say that $\frac{a}{b}, \frac{c}{d}$ *brackets* x.

Algorithm 1: Given $A = (0,0)$ has line-of-sight to $B = (q,p)$, construct $\mathcal{P}$ as an r-constrained path from A to B made up of line segments that have gradient $\frac{a}{b}$ or $\frac{c}{d}$ where $\frac{a}{b}, \frac{c}{d}$ is a Farey pair that brackets $\frac{p}{q}$.

$\mathcal{P} \leftarrow \lfloor AB \rfloor$
$\ell, u \leftarrow B$
for $m = 0 \dots pb - qa - 1$ **do**
    $u \leftarrow (\ell_x - d, \ell_y - c)$ ▹ *Move u to $\mathcal{L}_{m+1}$.*
    Remove from $\mathcal{P}$ the points that are strictly between u and ℓ
    $\ell \leftarrow$ Point in $\lfloor AB \rfloor \cap \mathcal{L}_{m+1}$ that is below AB with ℓ_x minimal ▹ *Move ℓ to $\mathcal{L}_{m+1}$.*
    Remove from $\mathcal{P}$ the points that are between ℓ and u and not on $\mathcal{L}_{m+1}$
return $\mathcal{P}$

Remark. Farey pairs have been used previously to approximate line segments; see the work of Harabor et al. (2016) and Rivera et al. (2017) for example. This article proves that the approximations exist and provides an explicit construction.

Algorithm 1 can be understood intuitively by considering a piece of elastic that is stretched from A to B. Put pins into the elastic wherever it crosses a cell edge. Slide the pins down to their closest respective cell corners. The elastic is then a path from A to B (Figure 3). Furthermore we still have a path if any of the interior pins are removed (that is, other than the first or last one) *provided* that whenever a pin is removed, the elastic restores upwards (Figure 4). Now consider the lines of gradient $\frac{a}{b}$ that pass through at least one cell corner. Restrict and enumerate the lines as $\mathcal{L}_m$ such that the first line passes through B and the last line passes through A (Figure 5). We can 'slide down' each $\mathcal{L}_m$, removing pins as we go, thereby building line segments of gradient $\frac{a}{b}$. We then 'step off' to $\mathcal{L}_{m+1}$ via a line segment of gradient $\frac{c}{d}$ (Figure 6). The elastic is then an r-constrained path from A to B made up of line segments that have gradient $\frac{a}{b}$ or $\frac{c}{d}$.

We need to prove that Algorithm 1 is correct. We show first that the 'stepping off' procedure is valid.

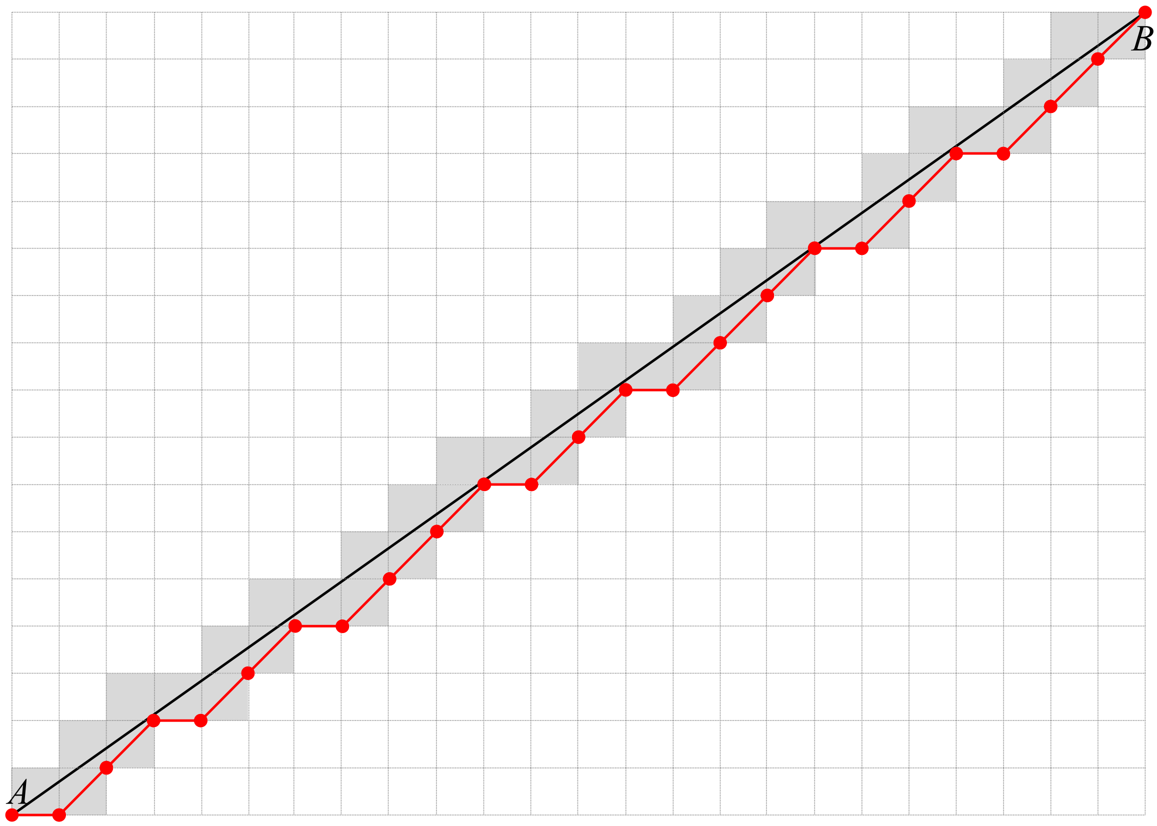

Figure 3: If A has line of sight to B then $\lfloor AB \rfloor$ is a path from A to B.

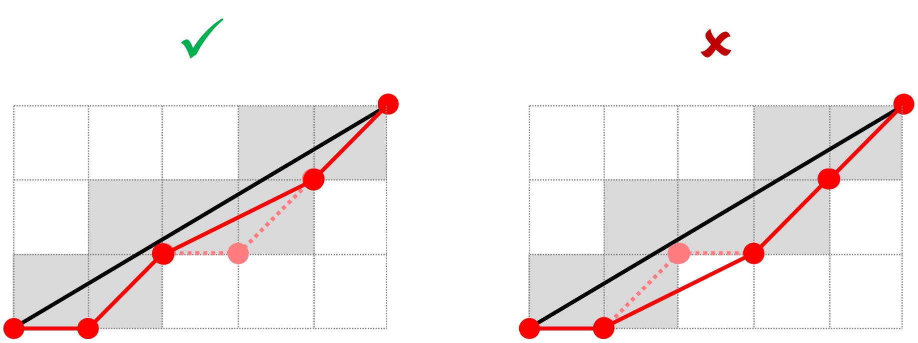

Figure 4: Removing pins yields a path provided the elastic restores upwards.

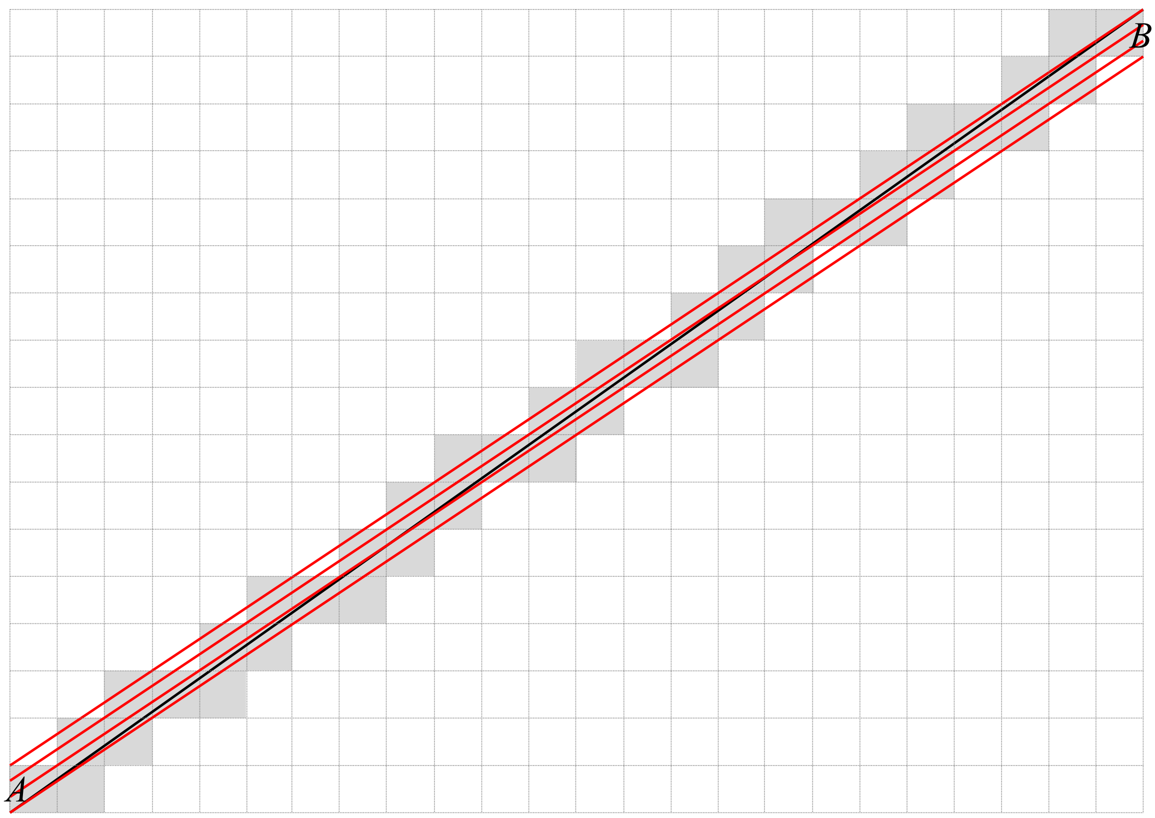

Figure 5: Lines of gradient $\frac{a}{b}$ that pass through at least one cell corner, denoted $\mathcal{L}_m$ so that the first line passes through B and the last line passes through A.

a)

b)

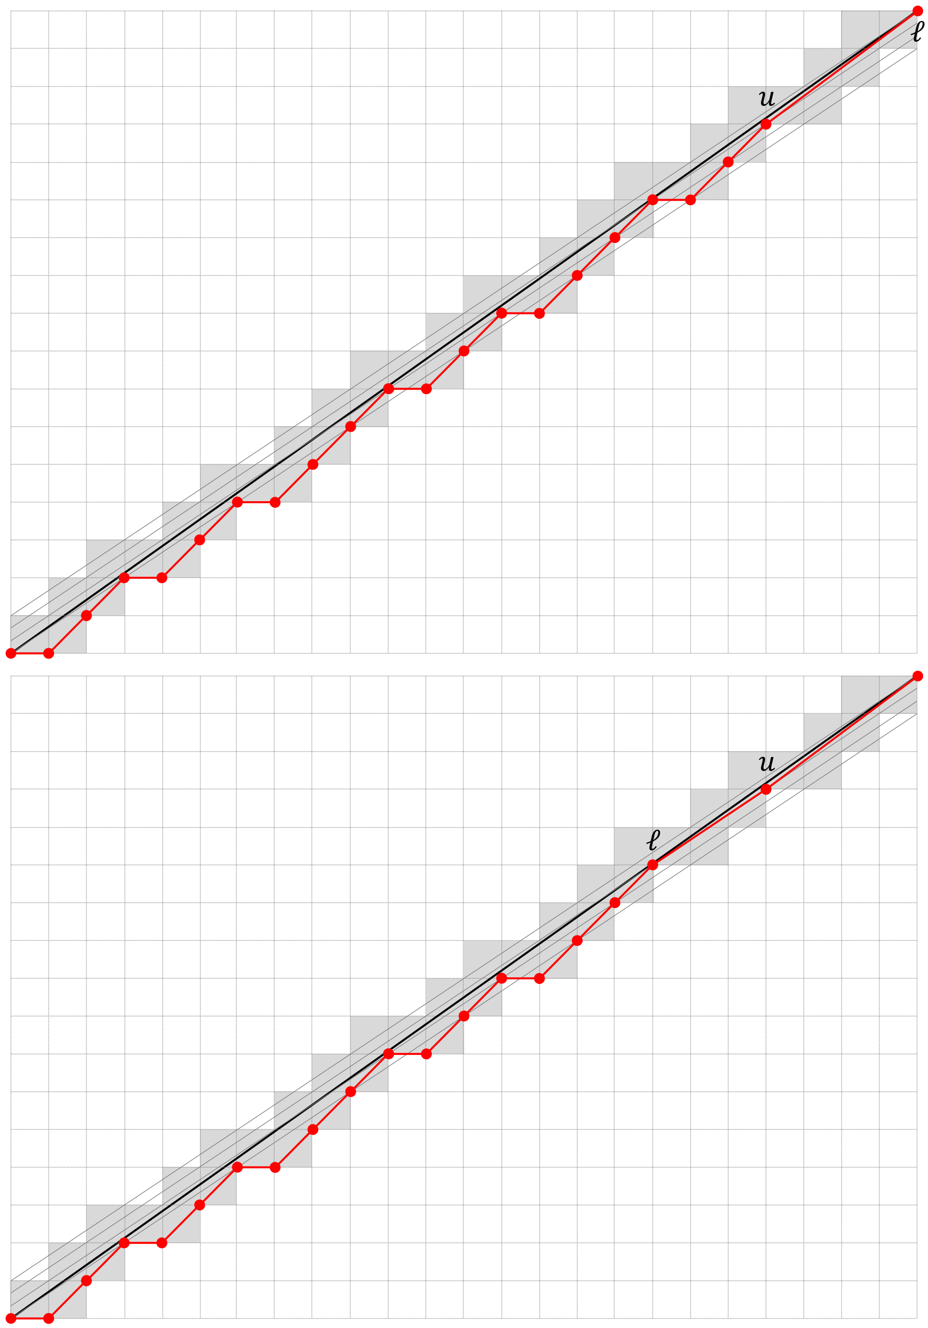

Figure 6: a) ‘Step off’ from $\mathcal{L}_1$ to $\mathcal{L}_2$ (lines 4–5). b) ‘Slide down’ $\mathcal{L}_2$ (lines 6–7).

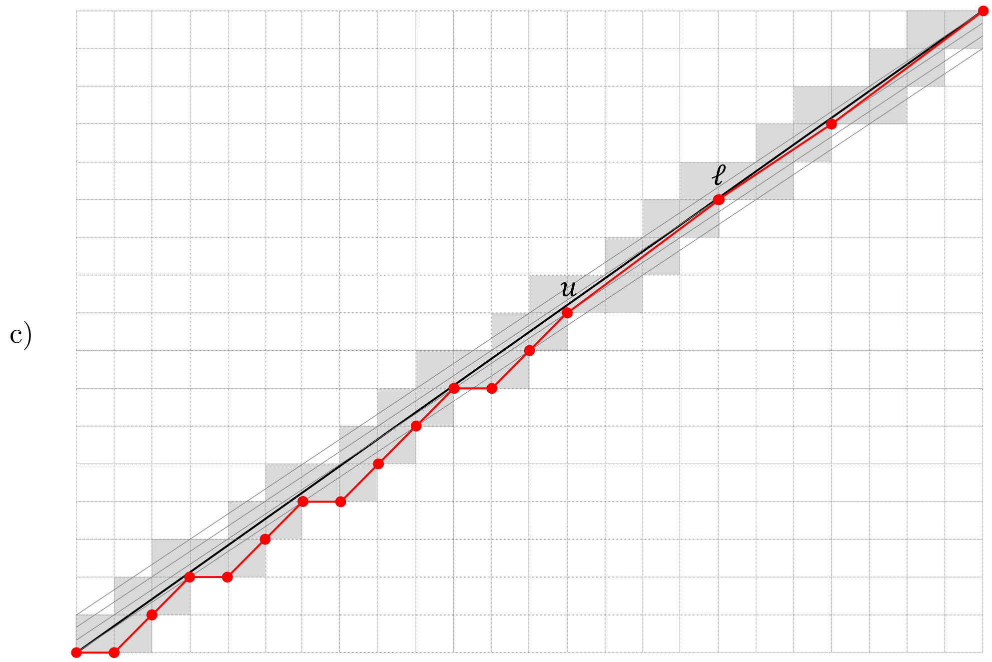

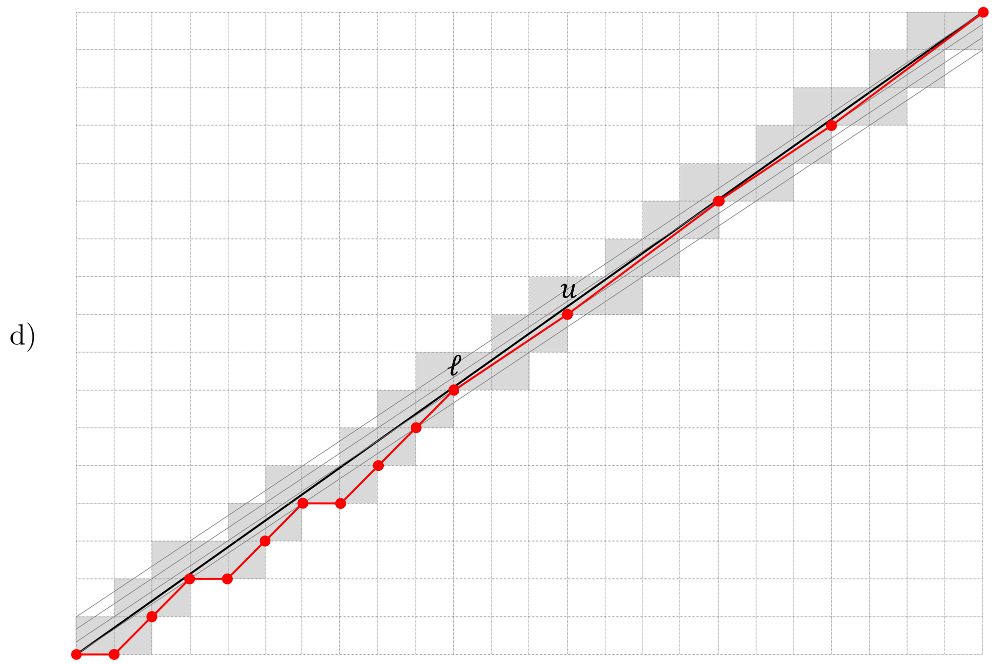

Figure 6: c) 'Step off' from $\mathcal{L}_2$ to $\mathcal{L}_3$ (lines 4–5). d) 'Slide down' $\mathcal{L}_3$ (lines 6–7).

Lemma 5 (The 'stepping off' is valid)**.** Suppose that ℓ is the point in $\lfloor AB \rfloor \cap \mathcal{L}_m$ that is below AB with ℓ_x minimal. Set $u = (\ell_x - d, \ell_y - c)$. Then $u \in \mathcal{L}_{m+1}$ and $u \in \lfloor AB \rfloor$.

Proof:

1. *Proof that* $u \in \mathcal{L}_{m+1}$*.* To satisfy $(\ell_x, \ell_y) \in \mathcal{L}_m$ we must have $\ell_y = \frac{a}{b}\left(\ell_x - \left(q + \frac{m}{a}\right)\right) + p$. Now $\frac{cb}{a} = d + \frac{1}{a}$ from $\frac{a}{b}, \frac{c}{d}$ being a Farey pair so

$$\begin{aligned} \ell_y - c &= \tfrac{a}{b}\left(\ell_x - \left(q + \tfrac{m}{a}\right)\right) + p - c \\ &= \tfrac{a}{b}\left(\ell_x - d - \left(q + \tfrac{m+1}{a}\right)\right) + p \end{aligned}$$

   Hence $(\ell_x - d, \ell_y - c)$ satisfies the equation for $\mathcal{L}_{m+1}$. □

2. *Proof that* $u \in \lfloor AB \rfloor$*.* Let $L = (\ell_x - d, \ell_y - c)$. Then L is below AB as (ℓ_x, ℓ_y) is below AB and $\frac{p}{q} < \frac{c}{d}$. Let $U = (\ell_x - d, \ell_y - (c - \frac{1}{b}))$. Now $(\ell_x - b, \ell_y - a)$ is above AB by the minimality condition so $(\ell_x - d, \ell_y - a - (d-b)\frac{a}{b})$ is also above AB from $\frac{a}{b} < \frac{p}{q}$. Moreover $-a - (d-b)\frac{a}{b} = -\frac{ad}{b} = -(c - \frac{1}{b})$ from $\frac{a}{b}, \frac{c}{d}$ being a Farey pair so U is above AB. Hence AB intersects LU. Finally L is on a cell corner (it has integer coordinates) and $\frac{1}{b} < 1$ so L is the cell corner immediately below the point where AB intersects the line $x = \ell_x - d$. Thus $L \in \lfloor AB \rfloor$. □

We now show that $\lfloor AB \rfloor$ is a path from A to B and that we still have a path if we are careful when removing points.

Definition.

1. Let $f, g : I \to \mathbb{R}^2$ be continuous functions from a closed interval I into the Euclidean plane $\mathbb{R}^2$ where $f(t) = (f_x(t), f_y(t))$, $g(t) = (g_x(t), g_y(t))$. We write $f \le g$ if $f_y(t) \le g_y(t)$ whenever $f_x(t) = g_x(t)$. Note that $\le$ is a partial ordering: $f \le f$ for all f, if $f \le g$ then $g \not\le f$, and if $f \le g$ and $g \le h$ then $f \le h$.

2. Let $\mathcal{Q} = \{s_k\}_{k=0}^n$ be a sequence. The *linear interpolation of* $\mathcal{Q}$ is the function $f : [0, n] \to \mathbb{R}^2$ defined as $f(t) = (t - \lfloor t \rfloor)s_{\lceil t \rceil} + (1 - (t - \lfloor t \rfloor))s_{\lfloor t \rfloor}$. Let $\mathcal{Q}, \mathcal{Q}'$ be sequences and f, g be their linear interpolations. We write $\mathcal{Q} \le \mathcal{Q}'$ if $f \le g$.

Lemma 6. $\lfloor AB \rfloor$ is a path from A to B. Moreover if X is a subsequence fom $0 \ldots q$ that contains both 0 and q, and $\lfloor AB \rfloor \le \lfloor AB \rfloor_X$ then $\lfloor AB \rfloor_X$ is also a path from A to B.

Proof:

1. *Proof that* $\lfloor AB \rfloor$ *is a path from* A *to* B*.* For every $x \in 0 \ldots q$, $(x, \frac{p}{q}x)$ has line-of-sight to $(x+1, \frac{p}{q}(x+1))$. Moreover $(x, \frac{p}{q}x)$ can be moved to $(x, \lfloor \frac{p}{q}x \rfloor)$ without breaking that line-of-sight and likewise $(x+1, \frac{p}{q}(x+1))$ can be moved to $(x+1, \lfloor \frac{p}{q}x + 1 \rfloor)$. □

2. *Proof that* $\lfloor AB \rfloor_X$ *is a path from* A *to* B*.* The sequence $\lfloor AB \rfloor_X$ includes both A and B by construction. Now the region between $\lfloor AB \rfloor$ and $\lfloor AB \rfloor_X$ is accessible and $\lfloor AB \rfloor_X$ stays inside that region. □

Lemma 7 (The removal of points is valid)**.** Let $\mathcal{Q}_m$ be the sequence that is held by $\mathcal{P}$ at the start of the mth iteration of Algorithm 1. Then for all m, $\mathcal{Q}_m$ is a path from A to B.

Proof: We apply induction to show that $\mathcal{Q}_0$ is a path from A to B and that if $\mathcal{Q}_m$ is a path from A to B then so is $\mathcal{Q}_{m+1}$.

1. *Proof that $\mathcal{Q}_0$ is a path from A to B.* $\mathcal{Q}_0 = \lfloor AB \rfloor$ and then apply Lemma 6. □

2. *Proof that if $\mathcal{Q}_m$ is a path from A to B then so is $\mathcal{Q}_{m+1}$.* $\mathcal{Q}_{m+1}$ is generated from by removing points from $\mathcal{Q}_m$ so by Lemma 6 it is sufficient to show that $\mathcal{Q}_m \leq \mathcal{Q}_{m+1}$. Suppose the algorithm is on its mth iteration. Put ℓ' as ℓ at the start of line 4. By the minimality condition on ℓ_x imposed at line 6, there are no points in lines $\mathcal{L}_0 \ldots \mathcal{L}_m$ to the left of ℓ'. Now on the mth iteration, the points that are removed are to the left of ℓ' so the points that are removed are from $\mathcal{L}_{m+1}$ onwards. Hence $\mathcal{Q}_m \leq \mathcal{Q}_{m+1}$. □

Having proved that Algorithm 1 is correct, we are now equipped to complete this section: if A has line-of-sight to B then there exists an r-constrained path $\mathcal{P}$ from A to B *and moreover* there is an upper bound on $\gamma(\mathcal{P})$.

Proposition 2 (Existence of r-constrained paths with bounded differences in gradient)**.** Let r be a positive integer. Suppose that A has line-of-sight to B and choose a grid-aligned coordinate system such that $A = (0, 0)$ and $B = (q', p')$ with $q' > 0$, $0 \leq p' \leq q'$. Then there exists an r-constrained path $\mathcal{P} = \{s_k\}_k$ from A to B and one of the following holds:

1. If $\frac{p'}{q'} = \frac{p}{q}$ for some $\frac{p}{q} \in \mathcal{F}_r$ then $\partial(s_k, s_{k+1}) = \frac{p}{q}$ for all k.

2. Otherwise let $\frac{a}{b}, \frac{c}{d}$ be the Farey pair from $\mathcal{F}_r$ that brackets $\frac{p}{q}$. Then for all k either $\partial(s_k, s_{k+1}) = \frac{a}{b}$ or $\partial(s_k, s_{k+1}) = \frac{c}{d}$.

Proof:

1. *Proof.* We partition AB into segments of the desired length: Let $K = \frac{q'}{q}$, set $x_k = kq$ for $k = 0 \ldots K$ and let $s_k = (x_k, \frac{p}{q} x_k)$. Then $0 < q \leq r$ and $0 \leq p \leq q$ so $\|s_{k+1} - s_k\|_\infty \leq r$. Furthermore $\partial(s_k, s_{k+1}) = \frac{p}{q}$ by construction. □

2. *Proof.* Use Algorithm 1 to construct $\mathcal{P}$ as a path from A to B with gradients $\frac{a}{b}$, $\frac{c}{d}$. □

2.3 Main Result

In previous sections, we have shown that:

1. *Proposition 1.* If $\mathcal{P}$ is a path from A to B then there is an upper bound on the inefficiency of $\mathcal{P}$ with respect to AB that can be calculated from $\gamma(\mathcal{P})$ (the difference in angle between the steepest and shallowest line segments of $\mathcal{P}$).

2. *Proposition 2.* If A has line-of-sight to B then there exists an r-constrained path $\mathcal{P}$ from A to B *and moreover* there is an upper bound on $\gamma(\mathcal{P})$.

We now combine these findings into our main result.

Theorem 1 (Existence of r-constrained paths with bounded inefficiency)**.** Let r be a positive integer and suppose that A has line-of-sight to B. Then there exists an r-constrained path $\mathcal{P}$ from A to B such that $\bar{\eta}(\mathcal{P}; AB) \leq \Delta$ where

$$\Delta = \begin{cases} 0 & \text{if } \partial(A,B) \in \mathcal{F}_r \\ 1 - \cos\left(\dfrac{\gamma}{2}\right) & \text{otherwise} \end{cases}$$

with $\gamma = \arctan\left(\frac{c}{d}\right) - \arctan\left(\frac{a}{b}\right)$ given $\frac{a}{b}, \frac{c}{d}$ is the Farey pair from $\mathcal{F}_r$ that brackets $\partial(A,B)$.

Proof. Choose a grid-aligned coordinate system such that $0 \leq \partial(A,B) \leq 1$. If $\partial(A,B) = 0$ or $\partial(A,B) = 1$ then $\mathcal{P}$ is constructed trivially. Otherwise construct $\mathcal{P}$ via Proposition 2 and obtain Δ from Proposition 1. □

By considering all of the Farey pairs from $\mathcal{F}_r$, we can obtain a worst-case inefficiency that relies solely on r.

Definition (Worst-case inefficiency)**.** Define

$$\hat{\Delta}(r) = 1 - \cos\left(\frac{\operatorname{arccot}(r)}{2}\right)$$

Lemma 8. For all Farey pairs $\frac{a}{b}, \frac{c}{d} \in \mathcal{F}_r$

$$1 - \cos\left(\frac{\gamma}{2}\right) \leq \hat{\Delta}(r)$$

where $\gamma = \arctan\left(\frac{c}{d}\right) - \arctan\left(\frac{a}{b}\right)$.

Proof. Observe:

1. *Let* $w = \max \frac{c}{d} - \frac{a}{b}$ *across all Farey pairs* $\frac{a}{b}, \frac{c}{d} \in \mathcal{F}_n$*. Then* $w = \frac{1}{n}$*, as attained by the pairs* $0, \frac{1}{n}$ *and* $\frac{n-1}{n}, 1$. We have $\frac{k}{n} \in \mathcal{F}_n$ for all $k = 0 \ldots n$ so $w \leq \frac{1}{n}$.

2. *Let* $f(v,w) = 1 - \cos\left(\frac{\gamma}{2}\right)$ *given* $\gamma = \arctan(v+w) - \arctan(v)$ *and* $0 \leq v < v+w \leq 1$*. Then* f *increases as* v *decreases and as* w *increases.* If $0 \leq v < v+w \leq 1$ then $\gamma > 0$ so $\sin(\gamma/2) > 0$. Moreover

$$\frac{\partial f}{\partial v} = \frac{1}{2}\sin\left(\frac{\gamma}{2}\right)\left(\frac{1}{1+(v+w)^2} - \frac{1}{1+v^2}\right) < 0$$

so f decreases as v increases. Likewise

$$\frac{\partial f}{\partial w} = \frac{1}{2}\sin\left(\frac{\gamma}{2}\right)\left(\frac{1}{1+(v+w)^2}\right) > 0$$

so f increases as w increases.

So write $\mathcal{F}_r = \{z_i\}_i$ and let $w = \frac{1}{r}$. By (1), $z_{i+1} - z_i \leq w$ for all i. Then by (2), $f(z_i, w)$ is maximized when $z_i = 0$; that is, when $\gamma = \arctan\left(\frac{1}{r}\right) - \arctan(0) = \operatorname{arccot}(r)$. □

We will need the following technical result, to generalize from line segments to paths.

Lemma 9. Let $\mathcal{Q} = \{s_k\}_{k=0}^n$ and $\mathcal{Q}'$ be the concatenation of $\{\mathcal{Q}'_k\}_{k=1}^n$ (with removal of duplicated points) where $\bar{\eta}(\mathcal{Q}'_k; s_{k-1}s_k) \leq c$ for all $k = 1 \dots n$. Then $\bar{\eta}(\mathcal{Q}'; \mathcal{Q}) \leq c$.

Proof.

$$\bar{\eta}(\mathcal{Q}'; \mathcal{Q}) = \frac{\sum_{k=1}^n \|\mathcal{Q}'_k\| - \sum_{k=1}^n \|s_{k-1}s_k\|}{\|\mathcal{Q}\|} = \sum_{k=1}^n \frac{\|\mathcal{Q}'_k\|}{\|\mathcal{Q}\|} \cdot \frac{\|\mathcal{Q}'_k\| - \|s_{k-1}s_k\|}{\|\mathcal{Q}'_k\|} \leq c \sum_{k=1}^n \frac{\|\mathcal{Q}'_k\|}{\|\mathcal{Q}\|} = c$$

□

In words: if the inefficiency of $\mathcal{Q}'_k$ with respect to $s_{k-1}s_k$ is bounded by some value c, for all k, then the inefficiency of the overall path is also bounded by c. We thus find that any path $\mathcal{P}$ has an r-constrained path $\mathcal{P}'$ with an inefficiency that we can write down.

Corollary 1. Let $\mathcal{P}$ be a vertex path from a start to a goal and r be a positive integer. Then there exists an r-constrained path $\mathcal{P}'$ from the same start to the same goal such that $\bar{\eta}(\mathcal{P}'; \mathcal{P}) \leq \hat{\Delta}(r)$.

Proof. Write $\mathcal{P} = \{s_k\}_k$. For each k, apply Theorem 1 to generate $\mathcal{P}'_k$ as an r-constrained path from s_k to s_{k+1}. Assemble $\mathcal{P}'$ as the concatenation of $\{\mathcal{P}'_k\}_k$ (with removal of duplicated points). Then $\mathcal{P}'$ is an r-constrained path with the same start and goal as $\mathcal{P}$. Moreover by Lemma 8 we have $\bar{\eta}(\mathcal{P}'_k; s_{k-1}s_k) \leq \hat{\Delta}(r)$ for each k. Result follows from Lemma 9. □

We check that our results apply to the *shortest* r-constrained paths.

Remark. Corollary 1 applies in particular if $\mathcal{P}$ is a shortest vertex path from a start to a goal (the corollary applies to any path). There still exists an an r-constrained path $\mathcal{P}'$ from the same start to the same goal such that $\bar{\eta}(\mathcal{P}'; \mathcal{P}) \leq \hat{\Delta}(r)$.

Lemma 10. Let $\mathcal{P}$ be a path from a start to a goal. If there exists path $\mathcal{P}'$ with the same start and goal such that $\bar{\eta}(\mathcal{P}'; \mathcal{P}) \leq c$ for some value c then $\bar{\eta}(\mathcal{P}''; \mathcal{P}) \leq c$ for all $\mathcal{P}''$ where $\|\mathcal{P}''\|_2 \leq \|\mathcal{P}'\|_2$.

Proof.

$$\bar{\eta}(\mathcal{P}''; \mathcal{P}) = \frac{\|\mathcal{P}''\|_2 - \|\mathcal{P}\|_2}{\|\mathcal{P}''\|_2} \leq \frac{\|\mathcal{P}'\|_2 - \|\mathcal{P}\|_2}{\|\mathcal{P}'\|_2} \leq c$$

where the second step uses Lemma 1. □

Corollary 2. Let $\mathcal{P}$ be a shortest vertex path from a start to a goal and r be a positive integer. Let $\mathcal{P}''$ be a shortest r-constrained path with the same start and goal as $\mathcal{P}$. Then $\bar{\eta}(\mathcal{P}''; \mathcal{P}) \leq \hat{\Delta}(r)$.

Proof. Corollary 1 constructs an r-constrained path $\mathcal{P}'$ with the same start and goal as $\mathcal{P}$ such that $\bar{\eta}(\mathcal{P}'; \mathcal{P}) \leq \hat{\Delta}(r)$. We have $\|\mathcal{P}''\|_2 \leq \|\mathcal{P}'\|_2$. Result follows from Lemma 10. □

We finish with some results of practical interest.

Remark. We have calculated Δ such that $\bar{\eta}(\mathcal{P}';\mathcal{P}) \leq \Delta$ for shortest vertex paths $\mathcal{P}$ and r-constrained paths $\mathcal{P}'$. Meanwhile Nash (2012, Theorem 1) found R such that $\|\mathcal{P}'\|_2 \leq R \cdot \|\mathcal{P}\|_2$ in the case of $r = 1$ (in a comprehensive study of grid paths on regular grids in two and three dimensions). We note that $\bar{\eta}(\mathcal{P}';\mathcal{P}) \leq \Delta$ if and only if

$$R = \frac{1}{1-\Delta}$$

Corollary 3 (Confirming Nash, 2012)**.** The shortest vertex paths from a start to a goal are at most 8 percent shorter than the shortest 8-connected paths.

Proof. Applying Corollary 2 with $r = 1$ yields $\bar{\eta}(\mathcal{P}';\mathcal{P}) \leq 8$ percent. Furthermore

$$\frac{1}{1-\hat{\Delta}(1)} = \frac{1}{1-\left(1-\cos\left(\frac{\operatorname{arccot}(1)}{2}\right)\right)} = \frac{2}{\sqrt{2+\sqrt{2}}}$$

which matches the value reported for R (Nash, 2012, Table 3.3). □

Corollary 4. The shortest vertex paths from a start to a goal are at most 1 percent shorter than the shortest 4-constrained ones.

Proof. Applying Corollary 2 with $r = 4$ yields $\bar{\eta}(\mathcal{P}';\mathcal{P}) \leq 1$ percent. □

3. Experiments

We have calculated upper bounds for the difference in length between shortest vertex paths versus r-constrained ones. We now compare our calculations with experiments.

3.1 Unblocked Terrain

We first consider an unblocked terrain, namely 113×113 cells that are all accessible. For $y = 0, 1, 2, ...113$, we let $A = (0, 0)$, $B = (113, y)$, and $\mathcal{P} = \{A, B\}$. We apply Dijkstra's algorithm to find a shortest r-constrained path $\mathcal{P}'$ from A to B.

Figure 7 charts the results for $r = 1 \ldots 6$. The charts use the angle from the x-axis to AB as the independent variable and inefficiency as the dependent variable. We plot the inefficiency as observed for r as a thin black line. The solid red line shows the bounds on inefficiency as predicted by Theorem 1. We see that these bounds are tight. The dashed blue line shows the worst-case bound predicted by Corollary 2.

3.2 Blocked Terrain

Blockages in the terrain will affect the existence or nature of vertex paths. But under our theory, if there exists a vertex path from a start to a goal then there exists an r-constrained path with an inefficiency that we can bound. To examine this prediction, we consider terrains of 113×113 cells in which a given percentage are accessible. We set A as an accessible location in the bottom-left quadrant and B as an accessible location in the right half. We apply A* to find $\mathcal{P}$ as a shortest vertex path from A to B and $\mathcal{P}'$ as a shortest r-constrained path.

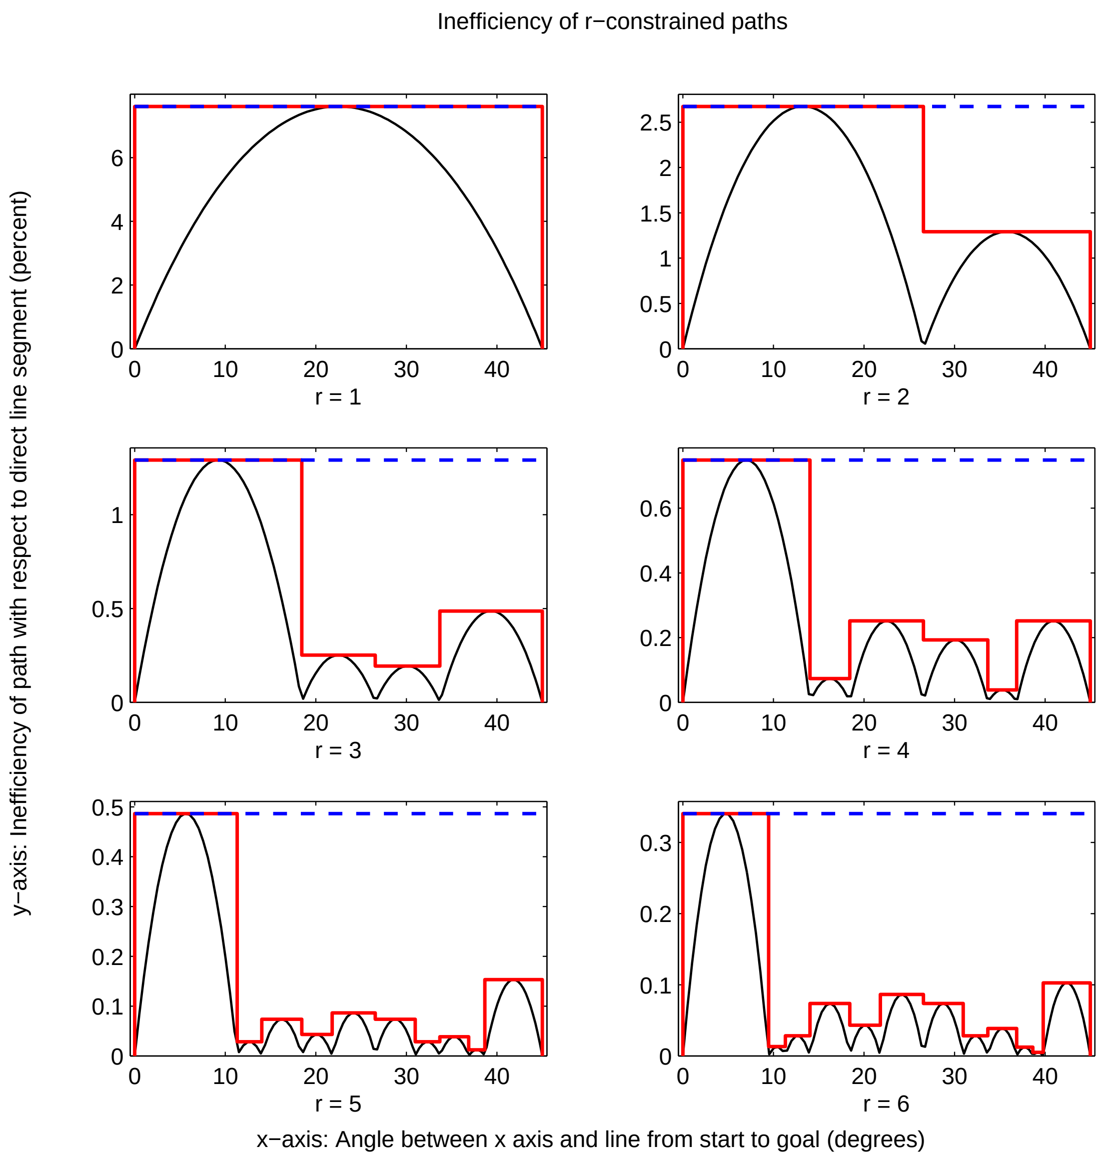

Figure 7: Inefficiency in unblocked terrain. For each chart, the solid red line shows the upper bound on inefficiency predicted by Theorem 1. The dashed blue line shows the worst-case bound predicted by Corollary 2. The thin black line plots the inefficiency as measured for the given r.

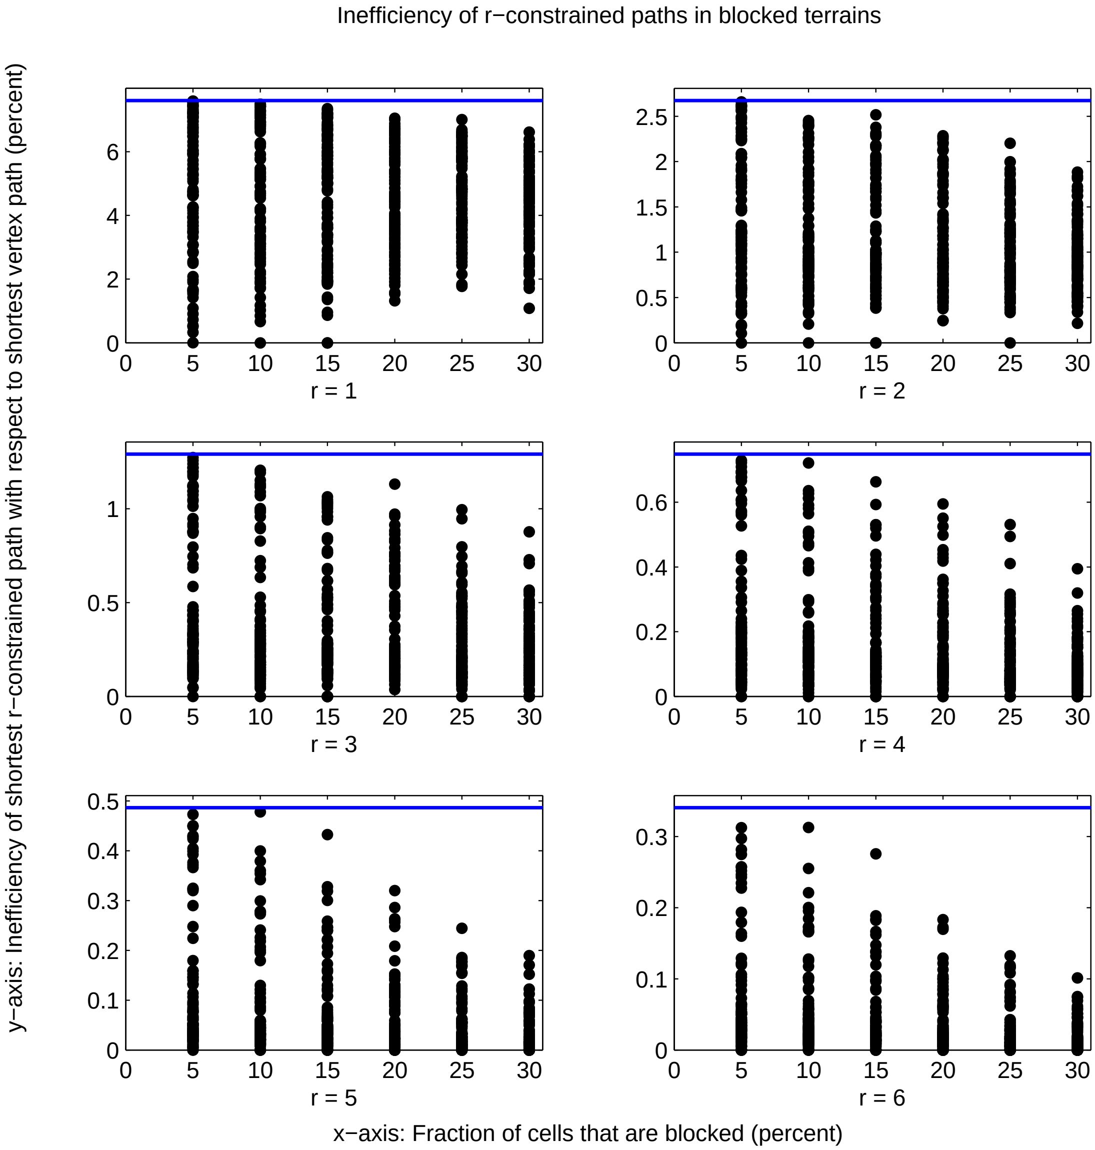

Figure 8: Inefficiency in blocked terrain. For each chart, the blue line shows the worst-case bound predicted by Corollary 2. Each dot shows the inefficiency recorded for a shortest r-constrained path with respect to a shortest vertex path.

Figure 8 charts the results for terrains where $5, 10, \ldots 30$ percent of the cells are blocked and $r = 1 \ldots 6$. Each dot shows the inefficiency recorded for a shortest r-constrained path with respect to a shortest vertex path; 100 trials were performed where each trial consisted of a terrain, start and goal as described above. The blue line shows the worst-case bound predicted by Corollary 2. We see that the empirical inefficiency is less than or equal to the worst-case bound that we predicted.

We also see that the empirical inefficiency tends to decrease as the blockages increase (the predicted worst-case bound is no longer tight). The behaviour can be accounted for as follows: as blockages increase, the opportunities for long, unimpeded line segments decrease. Hence the shortest vertex paths start to coincide with the shortest r-constrained ones and the inefficiency of the r-constrained paths becomes smaller.

4. Conclusion

The findings from this article can be used to trade path length for improved runtimes in a predictable manner. Suppose that we have a set of feasible locations and let $\mathcal{G}_r$ be the graph on those locations in which nodes are adjacent if they are within Chebyshev distance r. The paths through $\mathcal{G}_r$ have a 'digitization bias' that decreases as r increases. We quantified 'digitization bias' in terms of paths being inefficient and showed that the worst-case inefficiency is completely determined by r.

Now consider the runtime to search $\mathcal{G}_r$ deterministically for shortest paths; for example, by Dijkstra's algorithm or A*. There are $\mathcal{O}(r^2)$ nodes within Chebyshev radius r of a given node. Thus as r increases, the runtime to search $\mathcal{G}_r$ will increase as $\mathcal{O}(r^2)$. So increasing r will turn a 'good' vertex path into 'better' and eventually 'best-possible' but the additional runtime may be large. The same basic issue confronts all algorithms for short(est) vertex path planning.

Acknowledgments

The author appreciates the feedback and improvements from Edward Lo, Richard Taylor, Tansel Uras, Sven Koenig and the anonymous reviewers. This article is UNCLASSIFIED and approved for public release. Any opinions in this document are those of the author alone, and do not necessarily represent those of the Department of Defence (Australia).

References

Amante, C., & Eakins, B. (2009). ETOPO1 1 Arc-Minute Global Relief Model: Procedures, Data Sources and Analysis. NOAA Technical Memorandum NESDIS NGDC-24, National Geophysical Data Center, NOAA. Accessed 2015-09-25.

Bailey, J., Tovey, C., Uras, T., Koenig, S., & Nash, A. (2015). Path Planning on Grids: The Effect of Vertex Placement on Path Length. In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)*.

Harabor, D. D., Grastien, A., Öz, D., & Aksakalli, V. (2016). Optimal any-angle pathfinding in practice. *Journal of Artificial Intelligence Research*, *56*, 89–118.

Nash, A. (2012). *Any-angle path planning*. Ph.D. thesis, University of Southern California.

Nash, A., & Koenig, S. (2013). Any-angle path planning. *Artificial Intelligence Magazine*, *34*(4), 85–107.

Rivera, N., Hernández, C., & Baier, J. A. (2017). Grid Pathfinding on the 2^k Neighborhoods. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*.

Tsitsiklis, J. (1995). Efficient algorithms for globally optimal trajectories. *Automatic Control, IEEE Transactions on*, *40*(9), 1528–1538.

Uras, T., & Koenig, S. (2015). An Empirical Comparison of Any-Angle Path-Planning Algorithms. In *Proceedings of the Symposium on Combinatorial Search (SOCS)*, pp. 206–210.

Weisstein, E. W. (2017). Farey sequence. From MathWorld–A Wolfram Web Resource. http://mathworld.wolfram.com/FareySequence.html, Accessed 2017-05-25.