Rheinische Friedrich-Wilhelms-Universität Bonn
Mathematisch-Naturwissenschaftliche Fakultät

Theoretical Aspects of Intruder Search

MA-INF 1318 Manuscript Wintersemsester 2015/2016

Elmar Langetepe

Bonn, January 26th 2015

The manuscript will be successively extended during the lecture in the Wintersemester. Hints and comments for improvements can be given to Elmar Langetepe by E-Mail `elmar.langetepe@informatik.uni-bonn.de`. Thanks in advance!

Chapter 6

Escape Paths for the Intruder

In this chapter we would like to discuss a reverse situation. An intruder tries to escape from an environment as soon as possible. This work is inspired by the famous question of Bellman (brought up in 1956) who asked for the shortest escape path from an unknown forest.

As before we would like to consider geometric variants as well as more discrete situations. Again, the problem statement can be considered to be a game. The intruder has some abilities and tries to escape from the environment quickly whereas the adversary can manipulate the environment so that the intruder leaves the environment very late. We are looking for apropriate escape strategies for the intruder and consider different performance measures.

6.1 Lost in a forest

Assume that a simple region R in the plane is given which boundary is formally defined by a closed Jordan curve B that divides the plane in two simply connected regions. The intruder is located inside R and tries to find the boundary B as soon as possible by a deterministic escape path in the plane. We assume that the intruder has no sight system and only detects the boundary by touching it. The starting position p inside R and the rotation of R is unknown for the intruder but the exact shape of R is known. For example, somebody is located inside a dark forest of known shape and tries to get out of the forest as soon as possible. The deterministic escape path Π has to lead out of the region R for any starting position $p \in R$ and any rotation of R around p.

The performance measure for the path Π is simply its length. There will be a worst case starting position p of Π and a worst case rotation of R so that for p the full path length of Π is required to hit the boundary B. If no such point exists, there will be a better path Π of shorter length.

Considering such escape paths has a long tradition as mentioned above. Unfortunately, the optimal escape paths is only known for some very special shapes and totally unknown for general (polygonal) environments. We first discuss some simple convex situations where the diameter is optimal.

Conversely, the problem can also be considered as a covering problem. Consider the class C_L of rectifiable curves in the plane of some lenght L. Find an environment R of *small size* that can be rotated and translated so that it covers any curve of C_L. In the literature for $L = 1$ such covers are also denoted as *worm covers*. So we are searching for environments that are worm covers and but at least one worm finally touches the boundary from any starting situation.

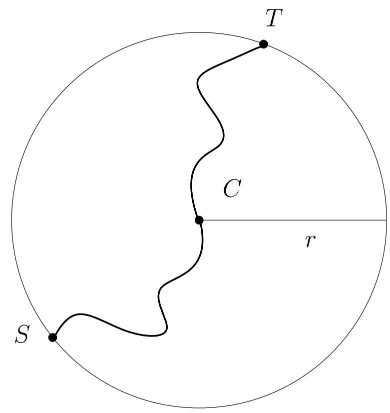

Figure 6.1: A cirlce of radius r covers any path of lenght $< 2r$. The diameter is a shortest escape path.

6.1.1 Simple examples and the diameter path

Let us assume that R is a circle of radius r.

Theorem 69 *The shortest escape path from a circle R of radius r is given by the diameter segment of R. Conversely, a circle of radius $\frac{1}{2}$ is a worm cover.*

Proof. Let us assume that the optimal escaps path Π of lenght L is given. We consider the point C on L that splits Π into two parts of lenght $L/2$. Let S denote the starting point of Π and T denote its endpoint. Now place C to the center of the circle; see Figure 6.1. In order to leave the circle at least S or T has to touch the boundary of R, thus $|CT| \geq r$ or $|CS| \geq r$. on the other hand the diameter d of R of lenght $2r$ is an escape path. □

Interestingly, also the semicircle has the same escape path and a semicirlce of diameter 1 is also a worm cover. The proof is a bit more complicated and was given by A. Meir and manifested by Wetzel (1973).

Theorem 70 *The shortest escape path of a semicircle R of radius r and diameter $2r$ is given by the diameter. Conversely, a the semicircle of radius $\frac{1}{2}$ is a worm cover.*

Proof. Let us assume that the escape path is a path with start end endpoint S and T. We rotate the path so that ST is in parallel with the base line B_l of the semicircle and we also translate the segment (and the path) so that there is a single tangent point I on the base B and all other points of the path Π lie above B_l as depicted in Figure 6.2. Consider an arbitrary point $X \in Pi$, w.l.o.g. we assume that X lies inside the path from S to I. We also consider the reflections of S' and T' of S and T along the base line B_l. The segment ST' intersects the base line at some point O. The length of the segment SX is shorter than Π_S^X. By refelection the length of the segment XT' is shorter than the path Π_X^T.

Now we translate the construction so that O is the center of the semicircle. We would like to argue that $|XO| \leq r/2$ holds. This means that any X is inside the semicircle which gives the conclusion.

We consider the triangle SXT' where O divides ST' into two parts of the same length. By geometry we know that the *median* XO is of the triangle is shorter that $\frac{1}{2}$ the length of the adjacent sites.

This means that

$$|XO| \leq \frac{1}{2}(|XS| + |XT'|) \leq \frac{1}{2}(\Pi_S^X + \Pi_X^T) < r\,.$$

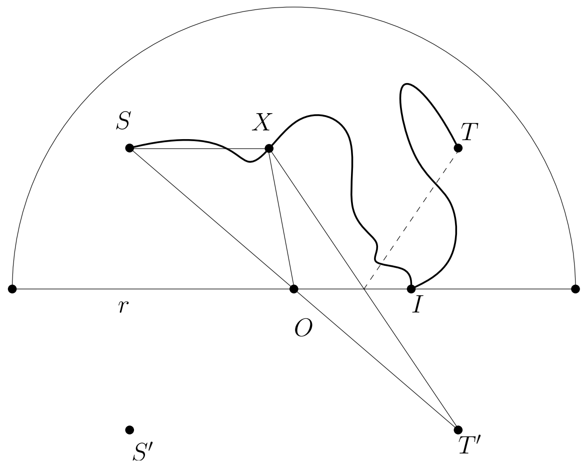

Figure 6.2: A semicirlce of radius r covers any path of lenght $< 2r$. The diameter is a shortest escape path.

□

Exercise 31 *Show that the* median *of the triangle is always shorter than the average (factor* $\frac{1}{2}$*) of the length of the adjacent sites.*

Exercise 32 *Show that the above Theorems also hold for closed paths. Note that we considered open paths in the proofs above.*

In general for a given region R the diameter d is defined to be the longest shortest path between two points in R. The corresponding points are always located on the boundary of R (otherwise there exist two points connected by a longer shortest path). For convex objects any diameter (path) is always an escape path, it need not be an optimal escape path as we will see in the next section. But for some convex objects R the diameter is indeed optimal as we will prove now. Some fatness condition is required.

Exercise 33 *Show that for non-convex polygonal objects the diameter need not be an escape path. Can you define a path for any simple polygon, that is always an escape path?*

First we consider the rhombus R of diameter L and angle $\theta = 60\circ$ as depicted in Figure 6.3. We would like to show that any escape path has lenght at least L. As already mentioned the diameter of lenght L is an escape path. The following proof stems from Poole and Gerriets (1973).

Theorem 71 *The optimal escape path for the rhombus* R_α *of diameter L and angle* $\alpha = 60°$ *is given by its diameter.*

Proof. As in the previous proofs we split an optimal escape path of some lenght $L' < L$ into two halfs of lenght $L'/2$ and consider the mean point C. We let C slide along the shorter diagonal BE as shown and rotate the path so that the path Π_S^C is tangent to AB and the path Π_C^T is tangent to BD. Such a rotational center for $C \in BE$ and orientation of R always exists. (Note that in an extreme case C could be located at B and only touches both AB and BD.) Let $X \in AB$ and $Y \in BD$ denote the corresponding tangent points of Π_S^C and Π_C^T, respectively.

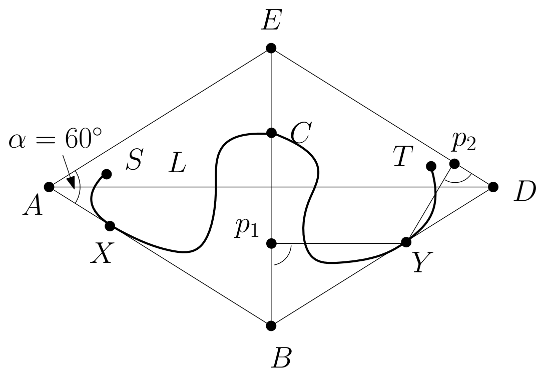

Figure 6.3: The rhombus R_α of diameter L and angle $\alpha = 60°$. The diameter is the shortest escape path since $|p_1Y| + |p_2Y|$ equals $L/2$ for any point Y on BD.

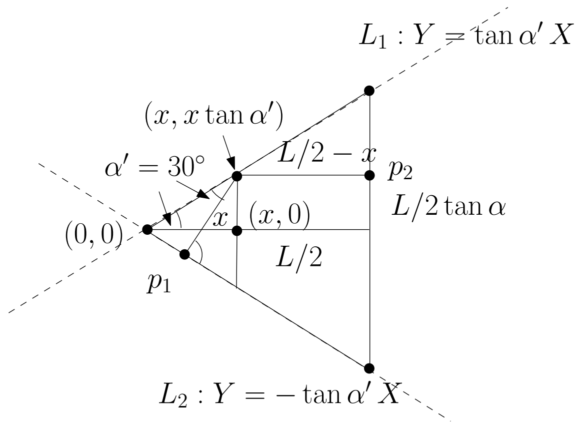

Figure 6.4: Parameterization for the conclusion in Theorem 71. The distance from $(x, x\tan\alpha')$ on L_1 to p_1 on L_2 also equals x for $\alpha' = 30°$.

Because L' is covered by R at least one path Π_S^C or Π_C^T has to hit the upper angle AED of R. W.lo.g. assume that Π_C^T hits AED. Consider the shortest path from Y to AED and to BE which meet ED at p_1 and BE at p_2 by angle $\pi/2$, respectively. This means that Π_C^T cannot be shorter than the sum of lenghts of p_1Y and p_2Y.

Finally, we show that $|p_1Y| + |p_2Y|$ equals $L/2$ for all $Y \in BD$, which concludes the proof. This holds by the geometric arguments shown in Figure 6.4. □

Exercise 34 *Is the shortest escape path always unique? Answer the question for convex and non-convex regions R.*

6.1.2 Besicovitch Zig-Zag path

Up to now all shortest escape paths where given by the diameter of the given convex object. This is not always true as we will show by a family of isosceles triangles T_α with base of length

$$b_\alpha = \sqrt{1 + \frac{1}{9\tan^2\alpha}}\,.$$

as shown in Figure 6.5 for $\alpha = 60°$ where T_α is a unilateral triangle. We will finally show that we can escape from the unilateral triangle of side-length 1 by a symmetric Zig-Zag path of lenght $\sqrt{\frac{27}{28}} < 1$, although 1 is obviously its diameter.

In general we consider α to be in the interval from roughly $52.24°$ up to $60°$. For $\alpha > 60°$ the base b_α is no longer the largest segment (and not the diameter any more), the reason for $\alpha \geq 52.24°$ is shown below. We show that for any T_α the shortest symmetric Zig-Zag-Path is an escape path for T_α with path lenght smaller than b_α. More precisely, the shortest symmetric Zig-Zag path will have lenght 1 and leaves T_α from any starting point.

The following result goes back to Coulton and Movshovich (2006). First, we define a symmetric Zig-Zag escape path. We orient T_α so that the base b_α runs in parallel to the X-axis and runs from $(0, 0)$ to $(b_\alpha, 0)$. The remaining segments l_α and r_α of T_α run above the X-asis in parallel along the lines $L_1 : Y = \tan \alpha X$ and $L_2 : Y = \tan \alpha(b - X)$, respectively as given in Figure 6.5.

The symmetric Zig-Zag consists of three consecutive segements of the same length and starts at the origing of T_α. Any segment has the same altitude h w.r.t. the base b_α. The last segment exactly touches the segment r of T_α. By construction any such path is an escape path for the corresponding T_α.

Now, we would like to construct a Zig-Zag path of length 1 for any T_α such that the path is the shortest symmetric Zig-Zag for a corresponding b_α. Such a path is the shortest, if its *straightened path* of the same length hits the line $L : Y = 3 \tan \alpha(b_\alpha - X)$ by a right angle as depicted in Figure 6.5 i). By congruence as shown in Figure 6.5 ii) we conclude that $\frac{1}{x} = \frac{b_\alpha}{1}$ which gives $x = \frac{1}{b_\alpha}$. Finally we determine b_α by $y = \tan \alpha \left(b_\alpha - \frac{1}{b_\alpha} \right)$ and $x = \frac{b}{1}$ and $x^3 + (3y)^2 = 1$ which gives

$$b_\alpha = \sqrt{1 + \frac{1}{9 \tan^2 \alpha}}.$$

For $\alpha = 60°$ we have $b_\alpha = \sqrt{\frac{28}{27}}$ and we can escape from the equilateral triangle of side-length 1 by a symmetric Zig-Zag path of lenght $\sqrt{\frac{27}{28}} < 1$, although 1 is obviously its diameter.

Exercise 35 *Verify the above formulas* $b_\alpha = \sqrt{1 + \frac{1}{9 \tan^2 \alpha}}$ *and* $b_\alpha = \sqrt{\frac{28}{27}}$ *for* $\alpha = 60°$.

For small α there might be other three-segment paths that also have distance 1 or even a shorter distance. This can happen for example, if a line $L_3 : Y = \tan(2\alpha)$ runs in parallel with L_2 as shown in Figure 6.6. This means $-3 \tan \alpha = \tan 2\alpha$ or $\tan \alpha = \sqrt{\frac{5}{3}}$. To avoid such situations we require $\alpha \geq \alpha_0$ where α_0 solves the equation.

Theorem 72 *For any* $\alpha \in [\arctan(\sqrt{\frac{5}{3}}), 60°]$ *there is a symmetric Zig-Zag path of lenght* 1 *that is an escape path of* T_α *smaller than the diameter* b_α.

Before we give a proof for the fact that the corresponding symmetric Zig-Zag escape paths are indeed optimal escape paths for any T_α we first introduce some other models and interpretations of the problem.

6.2 Different models and cost measures

Up to now we have considered the case that the intruder tries to escape from a geometric environment without any knowledge of its position. Let us assume that a bit more information is given and let us also consider a somewhat more discrete version. The intruder starts at the source of m long corridors, each of which finally lead out of the environment. The agent also knows the depth s_i but not the correspondance to the corridors.

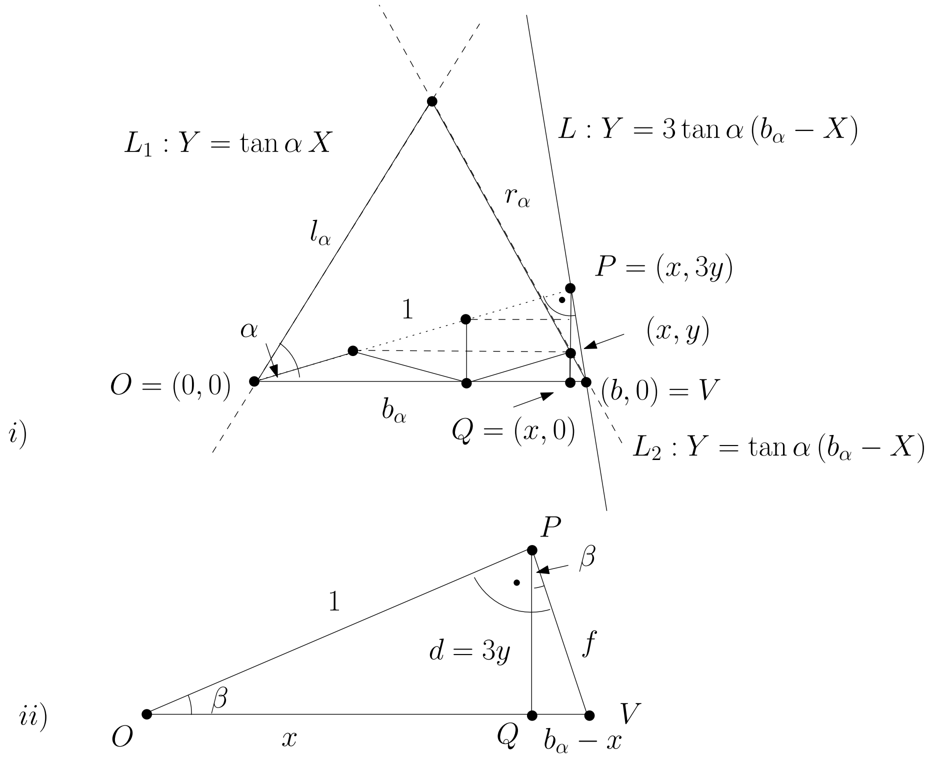

Figure 6.5: i) Any symmetric Zig-Zag path for T_α is an escape path. For the shortest such path the straightened segment mets the line $L : Y = 3\tan\alpha(b_\alpha - X)$ by a right angle. ii) Using the congruent triangles OPQ and OPV we have $\frac{1}{x} = \frac{b}{1}$.

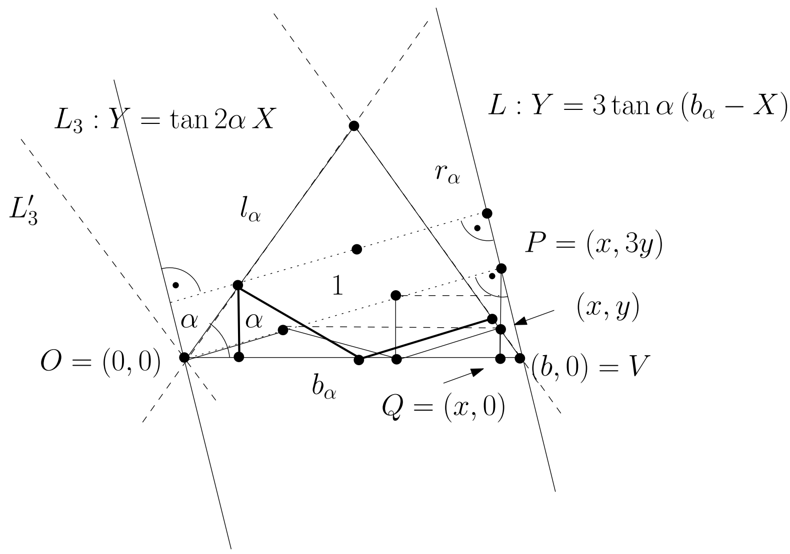

Figure 6.6: If α is too small, other Zig-Zag path might have a better performance (lenght ≤ 1).

Consider the situation where a set L_m of m line segments s_i of unknown length $|s_i|$ (which might represent dark corridors) are given and an agent has to find the end of only *one* arbitrary corridor as depicted in Figure 6.7(i). Just choosing a single corridor and move to its end might be very bad, if it is unfortunately the largest corridor. So in this escape problem the agent will move into one corridor s_{j_1} up to a certain distance x_1 and then check another corridor s_{j_2} for another distance x_2 and so on. Finally he will hit the end of one of the corridors with hopefully overall short path length.

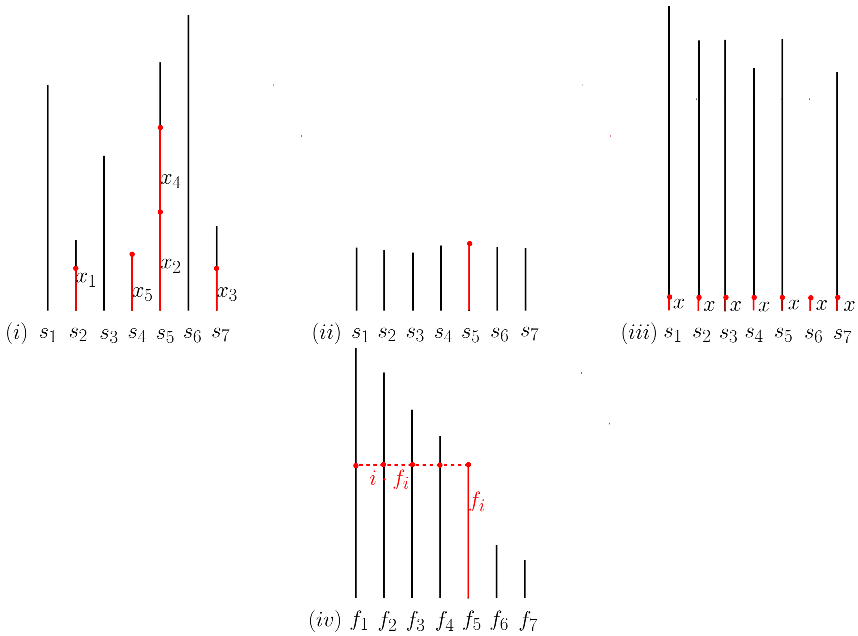

Figure 6.7: (i) Online searching for the end of one segment (or digging for oil) for $m = 7$ segments of unknown length as considered by Kirkpatrick. A reasonable strategy changes successively from one segment to the other with distances x_i on s_{j_i}. It is allowed to resume searching at the previous end-point with no extra cost (see x_2 and x_4 on $s_{j_2} = s_{j_4} = s_5$). The strategy reaches the end of segment s_5 after 5 movements $(x_1, \ldots, x_5)$. If at least the distance distribution is known, the problem is easier to solve. There are two extreme cases. (ii) If all segments have almost have the same length, it is reasonable to move along an arbitrary segment with largest distance, this is almost optimal. (iii) If there is one segment of very short length and all other segments are very long, one will find the end of the short segment by checking all segments with the shortest distance. This path is also short. (iv) In general the *discrete certificate* is defined for the given distance distribution. If $f_1 \geq f_2 \geq \cdots \geq f_m$ is the order of the length of the segments, it is always sufficient to check i arbitrary segments with length f_i and $\min_i i \cdot f_i$ is the best such strategy.

Kirkpatrick (2009) introduced this problem and also motivates the situation by the scenario of digging for oil at m locations s_i where the distance $|s_i|$ to the source of the oil of place s_i is not known. It is sufficient to get to the source of one place and the overall effort should be small. In this scenario it is allowed to resume the movement (or digging) for a location s_i at the endpoint where s_i was left at the previous visit; see Figure 6.7(i). There are no extra costs for moving to the previously reached depth at s_i.

Now we define a performance measure. Let us assume that all m distances $|s_i|$ are known but not the correspondence to the places s_i. In this case we can sort the m distances and obtain a discrete distance distribution of the length of the segments.

First, consider the extreme situations in Figure 6.7(ii) and (iii) If all segments s_i almost have the same length, a successful strategy will simply use the maximal length x among all segments, checks this for an arbitrary segment and will succeed with path length x in the worst case; Figure 6.7(ii). If on the other extreme (see Figure 6.7(iii)) the distance to a single source is small but the distances are very large to all others (and we only know this distribution), the best option is to check all segments successively by the small distance x. This gives an effort of at most $x \cdot m$ in the worst case when the small segment is found at the latest visit.

In general Kirkpartrick defines a *certificate* that takes the distance distribution into account. The segments are sorted by distance s_i. Let f_j with $j = 1, \ldots, m$ denote the sorted list of the lengths of the segments with decreasing distances f_j as shown in Figure 6.7(iv). Now for distance f_i there are exact i segments that are larger than or equal to f_i. Thus, after checking arbitrary i segments s_j with distance f_i it is clear that we will find the end of at least one segment. In the worst case the i largest segments have been checked including f_i. We summarize the above ideas in the following Theorem.

For a set L_m of m line segments s_i and a set of F_m of m length f_j, there is a permutation π so that $|s_i| = f_{\pi(i)}$ holds. Only L_m and F_m are given, the permutation is unknown. The agent can make use of a startegy as shown above. For any strategy A there will be a worst-case permutation $\pi(A)$ so that the traveling cost for finding at least one entrance is maximal.

For convenience it suffices to consider that only F_m is given and let us further assume that F_m is sorted, that is $f_1 \geq f_2 \geq \cdots \geq f_m$. Let $C(F_m, A)$ denote the maximum travel cost for algorithm A and list F_m which is attained for some permutation $\pi(A, F_m)$. For a given set F_m we are searching for the best strategy A and define the maximum-traversal-cost, maxTrav(F_m) by the minimum cost over all possible traversal strategies A, that is

$$\text{maxTrav}(F_m) := \min_A C(F_m, A)\,.$$

Theorem 73 *For a set of sorted distances F_m (i.e. $f_1 \geq f_2 \geq \cdots \geq f_m$) we have*

$$\mathit{maxTrav}(F_m) := \min_i i \cdot f_i\,.$$

Proof. Consider an arbitrary strategy A. If a strategy A moves less than $\min_i i \cdot f_i$ in total, the strategy has moved less that $j \cdot f_j$ for $j = 1, \ldots, m$. Thus, for any j, the strategy has visited less than j segments, say $k_j < j$ segments, up to distance f_j. Let us assume that the strategy has visiting depth $d_1 \geq d_2 \geq \cdots \geq d_m$ for the corridors. The adversary choose a permutation $\pi(A, F_m)$ so that the j-largest visiting depth d_j of A is applied to the segment of depth f_j. Thus the strategy has not checked the segment for f_1 up to the end, since $1 \cdot f_1$ has not been traveled in total. The segments of depths f_1 and f_2 has also not been visited up to the end since $2 \cdot f_2$ has not been travelled in total and the visiting depth at f_1 is as least as large as the visiting depth in f_2. Successively, the segments of depths $f_1, f_2, \ldots f_j$ has not been visited up to the end, since $j \cdot f_j$ has not been travelled and the visiting depths of $f_1, f_2, \ldots f_{j-1}$ are as least as large as the visiting depth of f_j. Altogether, the strategy is not successful for π with cost less than $\min_i i \cdot f_i$.

On the other hand if a strategy visits exactly i segments with depth f_i in the worst case the adversary presents the largest corridors but the ith-largest has length f_i. This means by $\min_i i \cdot f_i$ the strategy will be successful for any permutation. □

Now let us assume that in the online sense the distances are totally unknown to the agent. Kirkpatrick defines a general strategy that always approximates the certificate whithin a factor of $O(\text{maxTrav}(F_m) \log(\min(m, \text{maxTrav}(F_m)))$ for any list F_m of m totally unknown segments. It is shown that this factor is tight. The corresponding *dovetailling* strategy subdivides the overall digging length successively in a logarithmic way among an arbitrary order of the segments.

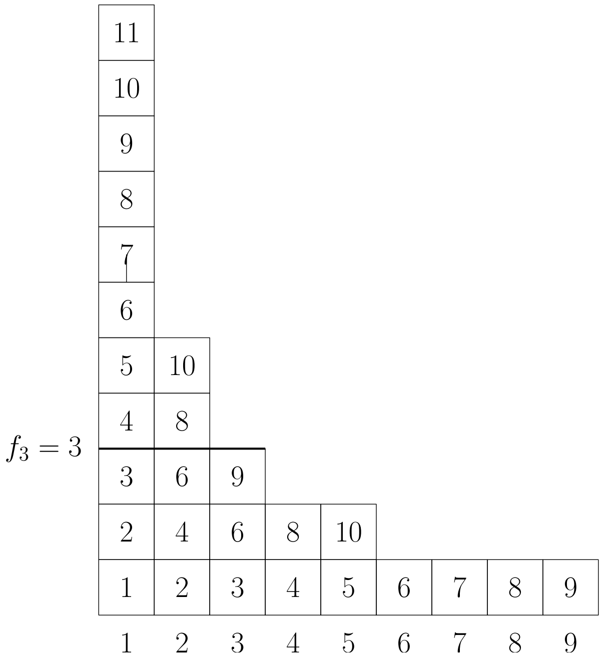

Figure 6.8: For $m = 9$ segments, in step $c = 1, 2, \ldots$ segment i is extended up to distance $\left\lfloor \frac{c}{i} \right\rfloor$ from left to right. For $f_3 = 3$, 3 segments will be covered with distance 3 during execution of step 9 on segment 3.

The online algorithm subdivides the cost w.r.t. a (discrete) hyperbolic function on the corresponding m segments. More, precisely the algorithm work in rounds $c = 1, 2, 3, 4, \ldots$ For any round c from left to right the path length on segment i is extended up to distance $\left\lfloor \frac{c}{i} \right\rfloor$. Figure 6.8 shows the extension scheme after 11 rounds for $m = 9$ segments. In any round any segment is extended at most by a distance of 1. If some segment i is extended in step c we have $\frac{c}{i} = t$ and the segments $1, 2, \ldots, i$ are all visited with depth at least t. This means that for any f_i in the interval from $t - 1$ to t, the strategy covers i segments up to distance f_i after step c for any i with $\frac{c}{i} = t$.

For convenience we assume that F_m contains only integer values. Now assume that for F_m we have $\min_i i \cdot f_i = j \cdot f_j$. There is some c with $c = j \cdot f_j$ and the strategy will be successful in this step. For this c the overall cost for the strategy is

$$\sum_{t=1}^{m} \left\lfloor \frac{c}{t} \right\rfloor \leq \sum_{t=1}^{\min(m,c)} \frac{c}{t} \leq c + \int_{1}^{\min(m,c)} \frac{c}{t} \, dt = c(1 + \ln \min(m, c)) .$$

Altogether, the following Theorem holds.

Theorem 74 *The hyperbolic traversal algorithm solves the multi-segment escape problem for any list F_m with maximum traversal cost bounded by*

$$D \cdot (\mathit{maxTrav}(F_m) \ln(\min(m, \mathit{maxTrav}(F_m)))$$

for some constant D.

Proof. By the considerations above for F_m with integer values for $D = 2$. □

Exercise 36 *Show that Theorem 71 also holds for non-integer F_m by doing the full analysis with a corresponding D.*

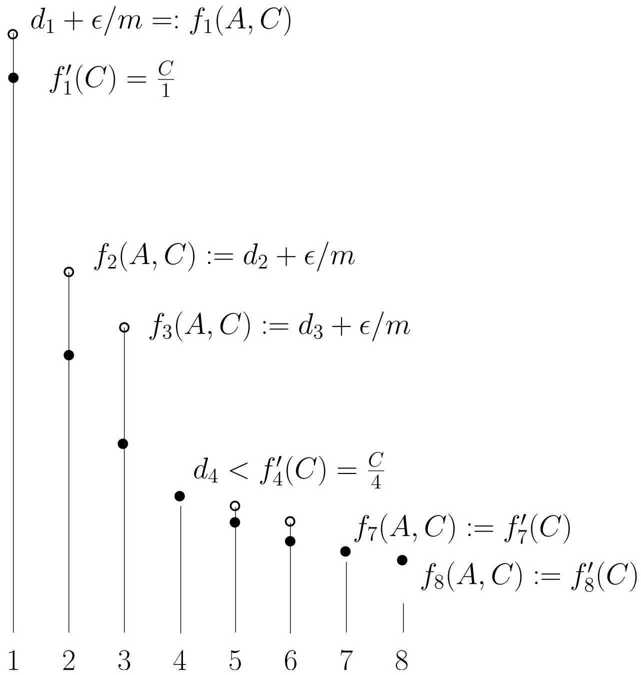

Figure 6.9: For $m = 8$ segments and some fixed C we define the lengths $f'_i(C) = \frac{C}{i}$ and let d be a constant so that $\sum_{i=1}^{8} f'_i(C) := d \cdot C \ln \min(C, m)$. A strategy A has visited the segments up to distance $d_1 \geq d_2 \geq \cdots \geq d_8$ and now attains the cost $\sum_{i=1}^{m} d_i \geq d \cdot C \ln \min(C, m)$ for the first time. For any segment j with $d_j < f'_j(C)$ we set $f_j(A,C) := f'_j(C)$. For any k with $d_k \geq f'_k(C)$ we set $f_k(A,C) := d_k + \epsilon/m$ for some arbitrarily small ϵ. There has to be some smallest index i so that $d_i < f'_i(C)$ and $i \cdot f'_i(C) = C$ is an upper bound of the optimal cost for $F_m(A,C)$.

Finally, we would like to argue that any deterministic strategy can be forced to have cost in the size of $d \cdot (\text{maxTrav}(F_m) \ln(\min(m, \text{maxTrav}(F_m)))$ and the corresponding dovetailing strategy is optimal in this sense. Therefore we assume that we consider overall cost C and define family of sets $F'_m(C)$ by $f'_i(C) = \frac{C}{i}$, so that $\text{maxTrav}(F_m(C)) = C$ and also $\sum_{i=1}^{m} f_i(C) := d \cdot C \ln \min(C, m)$ hold, by choosing d apropriately.

For large C we have situation as depicted in Figure 6.9. As long as the overall cost of a deterministic strategy A does not exceed cost $d \cdot C \ln \min(C, m)$ we do not let any corridors end.

For increasing visiting depth $d_1 \geq d_2 \geq \cdots \geq d_m$ of the strategy we consider the first moment in time where $\sum_{i=1}^{m} d_i \geq d \cdot C \ln \min(C, m)$ is attained. Immediately before not all corridors could have have been visited up to distance $f'_i(C)$, otherwise the overall path length $d \cdot C \ln \min(C, m)$ has been attained. Consider the smallest index i so that $d_i < f'_i(C)$. We finish the situation by fixing segment k with $d_k \leq f_k(C)$ to its final lenght $f_k(A,C) := f'_k(C)$ and fix all remaining segments with $d_j > f'_j(C)$ by $f_j(A,C) := d_j + \epsilon/m$ for some arbitrarily small ϵ. Note that $f_k(A,C), f_j(A,C) < f_i(A,C)$ holds for all corresponding $k, j > i$. The corresponding configuration can be denoted by $F_m(C, A)$. This means that $\text{maxTrav}(F_m(C, A)) \leq i \cdot f_i = \text{maxTrav}(F'_m(C)) = C$ holds and the lower bound is completed.

Theorem 75 *For any deterministic online strategy A that solves the multi-segment escape problem we can construct input sequences $F_m(A,C)$ so that A has cost at least $d \cdot C \ln \min(C, m)$ and maxTrav$(F_m(C, A)) \leq C$ holds for some constant d and arbitrarily large values C.*