

# Evolutionary learning of fire fighting strategies

Martin Kretschmer and Elmar Langetepe

University of Bonn, Department of Computer Science, D-53113 Bonn, Germany

**Abstract.** The dynamic problem of enclosing an expanding fire can be modelled by a simple discrete variant in a grid graph. While the fire expands to all neighbouring cells in any time step, the fire fighter is allowed to block  $c$  cells in the average outside the fire in the same time interval. It was shown that the success of the fire fighter is guaranteed for  $c > 1.5$  but no strategy can enclose the fire for  $c \leq 1.5$ . For achieving such a critical threshold the correctness (sometimes even optimality) of strategies and lower bounds have been shown by integer programming or by direct but often very sophisticated arguments. We investigate the problem whether it is possible to find or to approach such a threshold and/or optimal strategies by means of evolutionary algorithms, i.e., we just try to learn successful strategies for different constants  $c$  and have a look at the outcome. We investigate the variant of protecting a highway with still unknown threshold and found interesting strategic paradigms. **Keywords:** Dynamic environments, fire fighting, evolutionary strategies, threshold approximation

## 1 Introduction

In the field of motion planning, online algorithms or Computational Geometry (and of course in many other areas) there are many examples of (somewhat annoying) gaps between upper and lower bounds of interesting and important constants. For example in the field of online algorithms for the famous  $k$ -server problem in almost all metric spaces the best known lower bound on the competitive ratio is  $k$  whereas the best known upper bound is  $2k - 1$ , which gives a blind interval of  $[k, 2k - 1]$  for this value. The conjecture is that  $k$  is the tight bound; see for example [2]. Similarly the VC-dimension of  $L_2$ -visibility in simple polygons currently lies in the interval  $[6, 14]$ ; see for example [8]. The threshold has to be somewhere in between.

A challenging approach might be to close or reduce such gaps (or only get some more insight w.r.t. a tendency) by means of rather simple but efficient evolutionary or genetic approaches. In this paper for a suitable scenario in the context of motion planning in grid environments we would like to find out how far this general idea might work. We make use of an *Evolutionary Computation* approach and manipulate a population of solutions by natural selection and mutation such that a fitness gradually increases; see also [7,9,12,13]. Rather than analysing evolutionary algorithms theoretically as for example given in [1,3], we would like to analyse the power of such simple algorithms for getting insight in well-defined theoretical questions.

In this paper we concentrate on the context of discrete *fire fighting* in different theoretically motivated variants. An overview of theoretical results in this context is given by Finbow & MacGillivray [5]. Assume that in a grid-cell environment a cell that is on fire expands the fire from one cell to its four neighbouring cells in one time step. On the other hand the fire fighter can block some of the cells outside the fire in any time step. The number of cells that can be blocked is given by an asymptotic budget  $c \geq 1$  such that at any time step  $t$  we could have made use of  $\lfloor c \times t \rfloor$  blocked cells in total.

We examine two questions. It is well-known that for  $c > 1.5$  an expanding fire in the above model can be enclosed; see [11]. The result is obtained by a sophisticated recursive strategy idea. Optimality (minimum number of burned cells) can be obtained for example for  $c = 2$  by making use of ILP formulations; see [14]. This does not work well for smaller  $c$  because of the running times. On the other hand for  $c \leq 1.5$  no strategy can stop the fire, shown by a tricky proof in [4]. Therefore  $c = 1.5$  is the fixed threshold for this case.

For this well-understood scenario we make use of simple evolutionary rules and show that for  $c = 2$  we obtain the optimal strategy extremely fast. For  $c \geq 1.7$  we still obtain enclosurement results that seem to be close to the optimal. For  $c$  less than 1.6 our approach fails. The results are presented in Section 2.

The above first results might be seen as a test scenario for a new question considered in Section 3. For a protection budget  $c$  the task is rather than enclosing the fire, we would like to prevent a highway from being reached by the fire soon. Theoretical results and a fixed threshold for this setting are still unknown. We try to get an impression how reasonable strategies look like for different values of budget  $c < 1.5$ . It is more likely to make use of a single barrier close to the fire or is it recommendable to build (multiple) barriers away from the fire close to the highway? The focus here is that we get some ideas or insights by the use of evolutionary methods. In contrary to the former enclosurement problem we first make experiments and an ongoing task is to find formal proofs. The results and the corresponding conjectures are presented in Section 3.

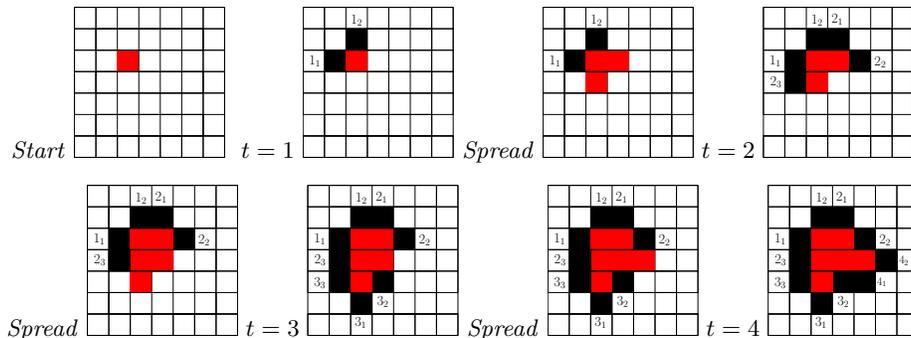
The main conclusion of our work is that simple, goal oriented evolutionary strategies could help to give insight into the solutions of dynamic motion planning problems. Especially, if such problems come along with a threshold question. The hope is that such approaches can also be used for similar problems.

## 2 Fire enclosurement in a discrete grid setting

Given an infinite grid graph with vertex set  $\mathbb{Z}^2$ . Each vertex represents a cell in a grid graph. In the following vertices and cells are handle as synonyms. The set of edges is given by  $\{((u, v), (x, y)) \mid |u - x| + |v - y| = 1\}$ , i.e. each cell is neighbour to the cell directly above, below, left, and right. A fire starts at  $(0, 0)$  and spreads over time. After each time step, all cells with a burning neighbour start burning as well.

In the first setting the goal is to enclose the fire, such that only a finite (minimal) number of cells is lost. To achieve this, a certain number of non burning cells can be protected at each time step, which will then never catch fire.

The number of cells that can be blocked is given by an asymptotic budget  $c \geq 1$  such that at any time step  $t$  we could have made use of  $\lfloor c \times t \rfloor$  blocked cells. A simple example for  $c = 2.7$  is shown in Figure 1. In the first step the fire fighter blocks  $\lfloor 1 \times 2.7 \rfloor = 2$  two cells outside the fire. After the fire spreads in the next step the fire fighter blocks  $\lfloor 2 \times 2.7 - 2 \rfloor = \lfloor 3.4 \rfloor = 3$  non-burning cells. Then the fire spreads again and in step 3 again  $\lfloor 3 \times 2.7 - 5 \rfloor = \lfloor 3.1 \rfloor = 3$  cells can be blocked by the fire fighter outside the fire. The fire spreads for the last time and by blocking  $\lfloor 4 \times 2.7 - 8 \rfloor = \lfloor 2.8 \rfloor = 2$  cells in the fourth step the fire is enclosed.



**Fig. 1.** An example for threshold  $c = 2.7$ . The fire starts at a single cell. At any time step, the fighter blocks the remaining cells of its overall budget  $\lfloor t \times 2.7 \rfloor$  outside the fire. Then the fire spreads. The protected (black) cells are labelled by time parameters. After 4 time steps the fire is enclosed.

It has been shown that a fire can always be enclosed protecting  $c = 2$  cells at each time step and it is impossible to do so with only one [6,14]. Finally, it was proved that a fire can always be enclosed when the average number  $c$  of protected cells exceeds 1.5 [11]. This bound is tight as shown by [4].

In the case of  $c = 2$  even an optimal solution (i.e. minimal number of burning cells) has been found by using Integer Linear Programming [14]. Compared to that, in the following we want to investigate how good a simple evolutionary inspired algorithm can solve this task and how close we can get to the thresholds. The first experiments also can be seen as a test scenario for the question of protecting a highway considered in Section 3.

## 2.1 A goal oriented evolution model

To use an evolutionary method, we require a formal description of a general strategy, which can be modified (mutation) and recombined (inheritance) to

obtain a new strategy. Additionally, we have to define a fitness function for the comparison of strategies. Intuitively (and also driven by the known theoretical results) it seems to be a good idea for a strategy to

- start close to the fire
- build a (more or less) connected chain of protected cells, trying to surround the fire

**Remark** We further confirm these intuitions by having tried other variants as well. Our evolutionary experiments showed, that strategies which start protecting vertices further away from the origin perform worse than strategies that start close to the origin, some results for this are presented in Table 1. Analogously, the experiments showed that multiple disconnected barriers (that finally might be connected) do not work well. We omit to show the corresponding experiments due to space constraints. We refer to Section 3 where we have similar results for the problem of protecting a highway. I.e., general disconnected genomes tend to run in a connected barrier construction. The following definition is designed to follow the above simple principles.

**Definition 1.** *A strategy consists of*

- a starting point  $(i, j)$
- a sequence of directions  $\{North, NorthEast, East, \dots\}$  and each direction is combined with the information whether to extend the front ( $F$ ) or the back ( $B$ ) of the chain

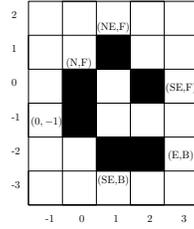
*For short the strategy is given by the starting point and a list of pairs  $(X, Y)$  with  $X \in \{N, NE, E, \dots\}$  and  $Y \in \{F, B\}$ .*

An example of a strategy (without a fire spread) is given in Figure 2. For the fixed starting cell  $(0, -1)$  the sequence  $((N, F), (NE, F), (SE, B), (SE, F), (E, B))$  is applied as follows. By  $(N, F)$  we extend  $(0, -1)$  *forward* by the cell  $(0, 0)$  in the north which now is the *new* front cell of the barrier. Then by  $(NE, F)$  relative to the new front cell we block the cell in direction north-east, which is cell  $(1, 1)$ . After that we apply  $(SE, B)$  for the current back end of the barrier which still is  $(0, -1)$ . The new back end cell is  $(1, -2)$  which lies south-east from  $(0, -1)$  and so on.

Notice, that such a strategy does not contain the information of the time at which the next vertex is protected. Instead, the next tuple of the sequence is applied, whenever we are allowed to protect an additional vertex.

The number of vertices that are protected per step is based on a bank account idea. We start with an initial budget and each time a vertex is protected, the budget decreases by 1. The budget has to remain positive but is always fully exhausted. After the fire has spread by one step, the budget increases by the fixed amount  $c$ . E.g.  $c = 2$  means we can protect exactly two vertices in any step. For  $c = 1.5$ , the number of protected vertices alternates between 1 and 2.

**Handling illegal genomes** The above genome design does not pay any attention to the restriction that we cannot protect already burning cells. To deal with



**Fig. 2.** Example of a strategy starting at  $(0, -1)$  with sequence  $((N, F), (NE, F), (SE, B), (SE, F), (E, B))$ . Each protected vertex is labelled by the tuple that caused its protection.

that we decide to use the following behaviour: Whenever the sequence tries to protect a cell that is already burning, we start a search for the next non-burning cell in clockwise or counter clockwise order, depending whether we want to extend the front or the back of the barrier, starting at the direction that is given by the sequence. For example in Figure 2, if  $(0, 0)$  is burning in the beginning, the application of  $(N, F)$  from  $(0, -1)$  results in blocking the cell  $(-1, 1)$ , which gives the new front.

**Fitness Evaluation** In order to determine the fitness of a strategy two values seem to be important. The time needed to enclose the fire and the total number of burning vertices. Since randomly initialized sequences will most likely not enclose the fire, we use the total number of burning vertices after a fixed simulation time  $t$ . This also gives rise to gradual improvements. For example in Figure 1 for a simulation time  $t = 3$  the given strategy has fitness 5, since 5 cells are burning at time  $t = 3$ . Note that we run arbitrary strategies with different simulation times (or steps).

## 2.2 Evolutionary Algorithm

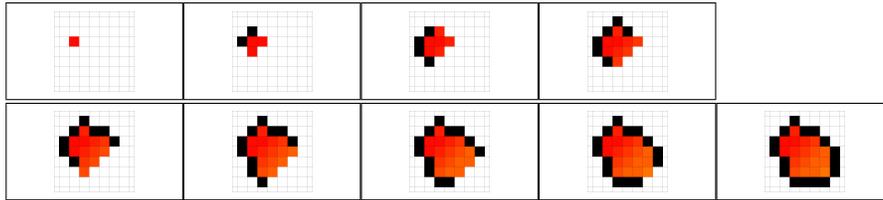
The following algorithm keeps improving a randomly initialized set of strategies until it is manually stopped. Besides the budget  $c$ , it has several parameters which determine its behaviour.

- Input:
  - $c$  budget income per time step
  - $n$  population size
  - $t$  number of simulation steps
  - $p$  mutation probability
  - $r$  ratio of parents kept after external selection
- Initialization: A population  $P$  of  $n$  randomly generated strategies (except the start point which is fixed to  $(0, 1)$ ), each strategy needs to have a sequence of length at least  $t \cdot c$

- Repeat
  - simulate any strategy of  $P$  for  $t$  time steps and determine its fitness
  - order  $P$  by fitness in increasing order and keep only the best  $\lfloor r \cdot n \rfloor$  strategies as parents
  - restock  $P$  again to size  $n$  by selecting two parent strategies and combining their sequences via single-point crossover
  - for each tuple in each sequence of  $P$ , change it with probability  $p$  to a new random direction and extension side (mutation)

Note that for speeding up the results of our simulation as presented in the next section for  $c < 2$  we decided to start the algorithm with an initial budget of 2. This allows us to protect two cell in the first step. Our experiments showed that this allows our algorithm to find successful strategies much faster and therefore also for smaller values of  $c$ . Asymptotically, there is no difference for the threshold. This small artefact might also be interpreted as a goal oriented approach.

### 2.3 Experimental results



**Fig. 3.** Example of a strategy found when protecting exactly  $c = 2$  vertices per step. Enclosed after 8 steps with 18 burning vertices.

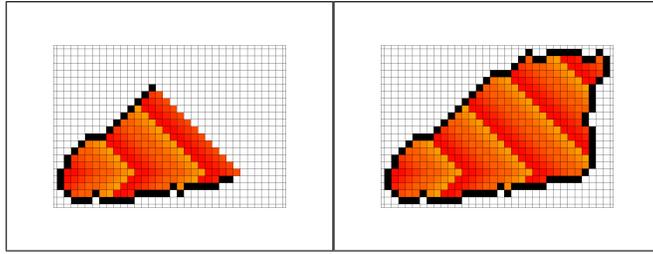
Figure 3 shows an optimal strategy that was found by evolution for the case  $c = 2$ . It takes 8 steps to enclose the fire and in the end 18 vertices are on fire. This is optimal for both time and number of burning vertices as shown in [6]. Surprisingly, it took only 84 generations in total until this strategy was found.

An example of a strategy that was found for  $c = 1.7$  after 1002 generations is depicted in Figure 4. A video of the successful strategy is shown in

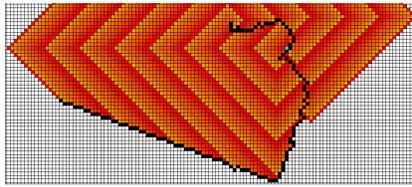
<http://tizian.informatik.uni-bonn.de/Video/1.7Enclosing.mp4>.

For even smaller values of  $c$ , our algorithm starts failing to find enclosing strategies. An example for  $c = 1.6$  is given in Figure 5. It seems that the strategy might be able to enclose the fire after a longer time, but even increasing the simulation time  $t$  did not lead to success.

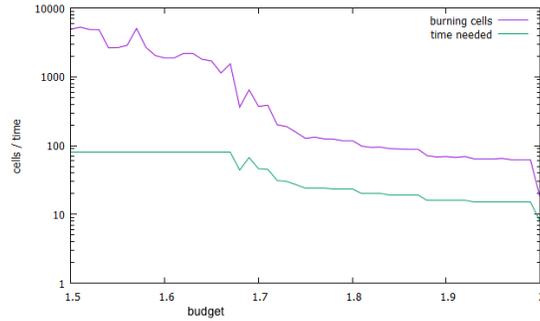
Figure 6 shows for which values of  $c$  we were able to find enclosing strategies. For constructing the figure we choose a simulation time of  $t = 80$ . As mentioned above increasing the simulation time did not help. One can see that for values



**Fig. 4.** Strategy found for  $c = 1.7$ . Enclosed after 46 steps with 371 burning vertices. The colored shading indicates how the fire spreads over time.



**Fig. 5.** Extract of a failing strategy for  $c = 1.6$ . Note that the fire expands on both sides of the barrier.



**Fig. 6.** An overview of the results. For convenience (results for different  $c$  were similar for any number of simulation steps) we used simulation time  $t = 80$ . For a given budget the green curve shows the time required for enclosing the fire and the blue curve shows the number of burning cells. We obtain positive (enclosure) results up to budget  $c$  slightly larger than 1.68. After that the green curve is just fixed to the restricted simulation time which indicates that the fire was not enclosed.

Start	Enclosing Time	Burning Vertices
(0, 1)	8	18
(0, 4)	23	156
(1, 3)	15	68
(2, 2)	24	161

**Table 1.** Fitness of best strategies found for  $c = 2$  and different starting points.

smaller than  $c \approx 1.68$  the building of the barrier continued until the simulation ended. This means that the fire was not enclosed.

So far, any strategy presented had a fixed start point neighbouring the origin of the fire. As an example we compare the strategy for  $c = 2$  mentioned above to strategies whose start point is fixed to a vertex four steps away from the origin. Up to symmetry there are three different coordinates for this.  $(0, 4)$ ,  $(1, 3)$  and  $(2, 2)$ . Table 1 shows the times required to enclose the fire using these different starting points, compared to the optimal strategy shown above. Starting further away from the fire takes longer to enclose the fire. We have similar results for other values of  $c$ .

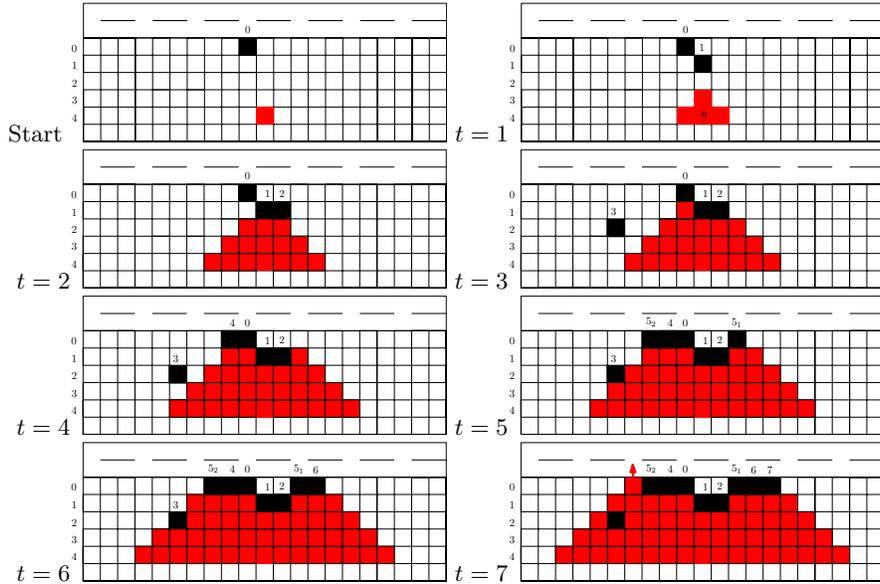
## 2.4 Fire encloement conclusion

At least for values of  $c$  a bit away from the overall tight threshold, the simple evolutionary goal oriented algorithm was able to find successful (and in the case of  $c \geq 2$  even optimal) strategies surprisingly fast. Successful strategies close to the threshold  $c = 1.5$  are not easy to find, even by the use of very general genomes and many simulation steps. This seems to be clear for the following reason. The corresponding successful connected barrier solution presented in [11] for any  $c = 1.5 + \epsilon$  makes use of four rounds. For any round the strategy behaves analogously but the next round starts with a 90 degree rotation. It is unlikely that a random approach will find the appropriate time for starting the next round and rotation.

## 3 Protection of a highway

Here we consider a different and new question. Conversely to the previous section we did not have any idea for a reasonable strategy and/or a threshold. The question is how long can we protect a highway (modeled by a line of cells) from the fire, if some budget  $c < 1.5$  is given. We would like to avoid that the fire touches a *line* very early? What is a reasonable strategy? Should we start close to the fire or close to the highway? Should we design a single connected barrier or more barriers which are partly disconnected?

In Figure 7 we give an example for a strategy for  $c = 1.2$ . This means that in the first 4 time steps the fire fighter makes use of a single blocking cell. In step  $t = 5$  the fire fighter can block two cells for the first time since  $\lfloor 5 \times 1.2 - 4 \rfloor = 2$  holds. Similar to the previous section we can also assume that in the start situation some constant cells are already blocked, this is indicated by the blocked cell of label 0 in Figure 7. Figure 7 has to be interpreted as follows. If  $c_{t-1}$  cells were used from the budget of the fire fighter after step  $t - 1$ , at the next time step  $t$ , the fighter first blocks  $\lfloor t \times c - c_{t-1} \rfloor$  cells outside the fire and then the fire spreads. After 7 time steps and the corresponding spread the fire reaches the highway. Note that the strategy stops in this moment.



**Fig. 7.** A fire fighter strategy for protecting a highway with budget  $c = 1.2$ . The fire starts at a single cell, one cell is initially protected. At any time step, the fighter first blocks the remaining cells of its overall  $\lfloor t \times c \rfloor$  budget outside the fire and then the fire spreads. After 7 time steps the fire reaches the highway.

### 3.1 Evolution models

Since the given problem was not theoretically analysed before, we first had to test several ideas experimentally in order to achieve a more goal oriented model. In contrary to the enclosurement scenario discussed before we do not know whether a connected strategy will lead to optimal or efficient solutions. So we first tried to allow general strategies that could protect arbitrary cells. We made use of a very simple coordinate based genome model, such that a strategy is simply defined by a set of cell coordinates defining which cells should be protected.

For such a set of cells, the cells are protected in their  $L_1$ -distance order from the origin of the fire, i.e., cells closer to the fire origin will be protected first. In the evolution process this behaviour forces that useless protections far away from the origin will be cancelled out more quickly. In principle the above principles allow us to define arbitrary strategies.

Altogether, we either make use of

- a connected genome (as in the previous section) or
- a coordinate genome described by a set of cell coordinates (as just mentioned)

### 3.2 Evolutionary Algorithm

We noticed that usually we do not get any improvements by the recombination of strategies. Therefore we changed the framework used in Section 2.2 and

restrict the algorithm to mutation only. This also means that we do not need to have a large population, instead we only initialize a single randomly generated strategy that will keep mutating. If a mutation leads to an improvement, the strategy keeps that mutation, otherwise it is undone. Altogether, this is a so-called (1+1)EA; see [12].

This process of improving a single strategy can easily be parallelized such that a larger set of single strategies keeps improving over time. This is very beneficial, because the final result often depends on the initialization and not every run leads to the best result. Another difference to the previous section is the fitness evaluation, which obviously has to be adjusted with respect to the problem definition.

**Fitness evaluation** For the enclosurement problem considered in Section 2.1 we tried to minimize the total number of burning cells. In this case we have used exactly this number for determining the fitness. Now we want to maximize the time the fire requires to reach the highway. It turns out that increasing this time value directly by a random mutation or recombination is very unlikely. Therefore we require a fitness evaluation that also allows for smaller and gradual improvements. To attain this we take into account how many vertices are burning and also their corresponding distance to the highway. A formal definition is given below.

For letting the algorithm run, actually we only need to be able to compare strategies pairwise. Fortunately, this can also be realized by our fitness function.

**Definition 2.** *Let  $S$  be a protection strategy for a given highway. By  $r(S)$  we denote the first moment in time when the fire reaches the highway, if  $S$  is applied. By  $d(S)_i$  we denote the number of burning cells with distance  $i$  to the highway after  $r(S)$  simulation steps.*

*Strategy  $S_1$  has a larger fitness than  $S_2$  if  $r(S_1) > r(S_2)$  holds or for  $r(S_1) = r(S_2)$  if  $d(S_1)_i < d(S_2)_i$  holds for the smallest index  $i$  where  $d(S_1)_i \neq d(S_2)_i$ .*

For example in Figure 7 the given strategy  $S$  has value  $r(S) = 7$ . We also have  $d(S)_0 = 1$ ,  $d(S)_1 = 9$  and  $d(S)_2 = 12$  and so on. So another strategy  $S'$  would have larger fitness, if  $r(S') > 7$  holds or for  $r(S') = 7 = r(S)$ , if we have for example  $d(S')_0 = 1$ ,  $d(S')_1 = 9$  and  $d(S')_2 = 11 < 12 = d(S)_2$ .

The main idea is that by trying to keep the fire farther away from the highway, finally also the overall time where the fire reaches the highway can be increased.

### 3.3 Experimental results

Similar to the enclosurement problem our implementation allows us to set or manipulate many different parameters and options for a goal oriented evolutionary process, such as the budget  $c$ , the strategy design (general genome or connected barriers), the population size (number of strategies optimized in parallel), the mutation rate, the fire source (distance to the highway), starting positions (for connected barriers), optional initial budget and so on.

**Videos** Finally and interestingly we mainly found two different strategic behaviours depending on the corresponding genomes, they will be explained precisely below. For convenience for  $c = 1.2$  we prepared two animations that show the finally attained best strategies for

**1. General genomes** Symmetric and alternating strategy:

<http://tizian.informatik.uni-bonn.de/Video/1.2SymAlt.mp4>

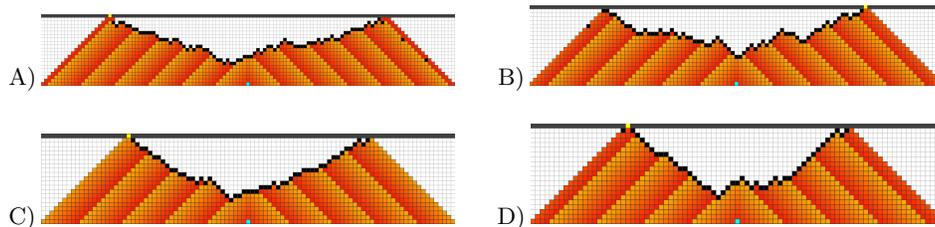
**2. Connected barriers** Asymmetric and diagonal strategy:

<http://tizian.informatik.uni-bonn.de/Video/1.2AsymDiag.mp4>

where in the second case of connected barriers sometimes also symmetric and alternating strategies were attained under circumstances explained below. The above strategies have been found after 156925 (1.) and 34226 (2.) generations.

**1. General genomes** First, we found out that the use of general genomes always (for different settings) mutate toward connected barriers; Figure 8 shows some of the finally attained strategies. All strategies show a similar behaviour. They start somewhere between the origin of the fire and the highway, usually a bit closer to the fire. Then any strategy continues to protect cells alternating between left and right, trying to keep the fire as long and as far away from the highway as possible. In the following we refer to such strategies as *symmetric and alternating*.

Note that any of the given strategies can be reconstructed such that the symmetric and alternating process is performed directly at the highway. The time where the fire reaches the boundary will not change in this case. Our fitness function simply prefers to shift the fire away from the highway.

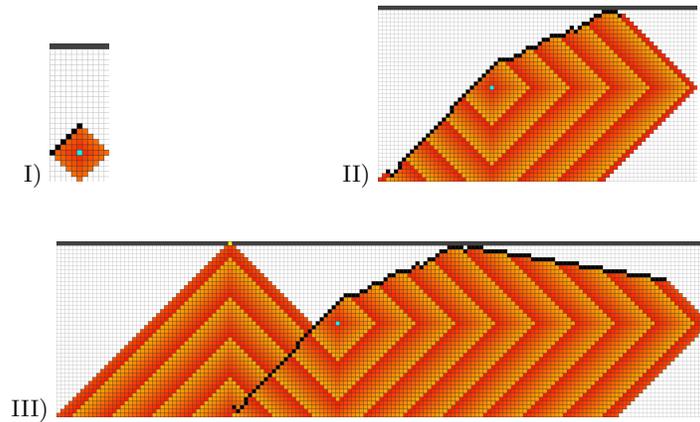


**Fig. 8.** Resulting strategies using the general coordinate genome for different values of the budget: A)  $c = 1.4$ , B)  $c = 1.3$ , C)  $c = 1.2$  D)  $c = 1.1$ . The fire starts 20 steps away from the highway. The highway was reached after 61, 54, 48 and 43 steps, respectively.

**2. Connected barriers** After that we again considered connected barriers with different starting positions below the origin. Depending on the distance between the start and the fire, we observed two different strategic behaviours which can be categorized as follows.

If the starting position is somehow chosen *too close* to the fire origin or *too close* to the highway we obtain strategies that behave in a symmetric and

alternating way as before. On the other hand if we somehow start at the *right* distance, the attained strategies suddenly performed different and a lot better. An example of such a strategy for  $c = 1.2$  is given in Figure 9, the behaviour of the strategy can be subdivided into three different phases which will be explained below. In contrast to the symmetric and alternating strategy which only kept the fire for 48 steps away from the highway with the same budget, the alternative strategy increases this time to 92!



**Fig. 9.** Three phases of a connected strategy protecting the highway for  $c = 1.2$ . I) Reach the level of the starting position of the fire diagonally from the center. II) Extend this diagonal, but simultaneously also try to shift the fire away from the highway on the other side until the fire reaches the highway there. III) Use the full budget to keep the fire away at this side, the fire *runs* to the highway at the other side. The strategy starts 5 steps above the origin. The fire reaches the highway after 92 steps.

In general these strategies can be subdivided into three phases. Figure 9 shows the end of each phase.

- I) Protect a diagonal downwards until a cell at the same level as the origin is reached. Starting  $n$  cells above the origin, this requires the protection of  $n + 1$  cells and this needs to be done before the  $n$ -th time step, because otherwise the fire would reach that cell first. This in turn requires  $c$  to be large enough. Or the other way round, given  $1.0 < c < 1.5$ ,  $n$  needs to be large enough such that  $n + 1$  cells can be protected after  $n$  steps  $\Leftrightarrow cn \geq n + 1 \Leftrightarrow n \geq \frac{1}{c-1}$ .
- II) Continue the diagonal downwards by one cell in every second step. Use the rest of the budget to keep the fire at the other end of the barrier as far away from the highway as possible. This procedure ends when the fire gets close to the highway.
- III) In order to protect the highway, from now on we are forced to protect at least one cell per step at the end close to the highway. Since protecting one

cell at every step at one end and one cell at every second step at the other end would require a budget of  $c \geq 1.5$ , the diagonal part of the barrier will be overrun by the fire making it impossible to continue this end at all. So the strategy will simply continue to hold the fire back at the upper part of the barrier until the fire reaches the highway on the other side. Again, because the fitness evaluation prefers fewer burning vertices close to the highway, the slope of the part built in this phase occurs.

We will refer to this behaviour as an *asymmetric and diagonal* strategy. Notice that for  $c = 1.5$  this leads immediately to a strategy that protects the highway infinitely. Furthermore this strategy can only be applied if the fire starts far enough away from the highway. The closer the budget  $c$  gets to 1.0, the more distance is required. If this distance is not available, there seems to be no better strategy than the symmetric and alternating one.

### 3.4 Highway protection conclusion

Using the evolutionary algorithm, we gained helpful insights into the highway protection problem. Both the symmetric and alternating and the asymmetric and diagonal strategy are promising candidates for optimal solutions. The choice between the two alternatives seem to depend on the possibility of building the diagonal of phase I). We found out that with one additional initial budget, the connected genome always run into the asymmetric and diagonal variant. By protecting two cells in the beginning, we can immediately finish the first phase by starting directly above and to the left of the fire. Without an initial budget the strategy has to fight for reaching the starting level of the fire source from the left. This can only happen if in comparison to the budget, the source lies sufficiently far away from the highway.

Considering phase III), there seems to be some room for a further recursive improvement. When the fire has overcome the diagonal part of the barrier in phase II) it will take the direct way to the highway. For a while we shift the fire away from the other side by using the full budget. The barrier is build with a given slope; see Figure 9 III). But this part could have been build also with budget 1 along the highway. Therefore the remaining budget can be used to protect the highway at the left hand side. Therefore we can consider the situation with a budget  $c' = c - 1$ . We found out that in this case the best strategy builds a symmetric and alternating barrier directly at the highway.

For values  $c \approx 1.5$  we never observed a strategy that was able to protect the highway infinitely long. We think that 1.5 is the threshold.

**Conjecture** For  $c < 1.5$  there is no strategy that protects an arbitrary highway from the fire. The best protection strategy either builds a single connected barrier symmetrically and alternating close to the highway or first the asymmetric and diagonal connected barrier strategy is applied. This depends on the relationship between the distance of the fire source to the highway and the given budget.

## 4 Future work on theoretical threshold questions

Besides proving and analysing the above conjecture theoretically, we finally would like to mention that there are other interesting upper and lower bounds which we analogously would like to attack by a goal oriented evolutionary approach. For example, similar to the subjects presented here there are other scenarios in discrete and continuous fire fighting settings that come along with a threshold. An interesting overview for such gaps is given in the CG Column by Klein and Langetepe[10]. Alternatively, one might also think of the protection for different objects, also formalized by a set of cells. Additionally, among many others, the two blind intervals (VC-dimension,  $k$ -server conjecture) mentioned in the very beginning are also worth considering.

**Acknowledgements** The authors would like to thank the anonymous reviewers for their valuable comments and suggestions.

## References

1. H.-G. Beyer, H.-P. Schwefel, and I. Wegener. How to analyse evolutionary algorithms. *Theoretical Computer Science*, 287(1):101 – 130, 2002.
2. M. Chrobak and L. L. Larmore. An optimal on-line algorithm for  $k$ -servers on trees. *SIAM J. Comput.*, 20(1):144–148, Feb. 1991.
3. S. Droste, T. Jansen, and I. Wegener. On the analysis of the (1+1) evolutionary algorithm. *Theoretical Computer Science*, 276(1):51 – 81, 2002.
4. O. N. Feldheim and R. Hod. 3/2 firefighters are not enough. *Discrete Applied Mathematics*, 161(1-2):301 – 306, 2013.
5. S. Finbow and G. MacGillivray. The firefighter problem: a survey of results, directions and questions. *Australasian Journal of Combinatorics*, 43(57-77):6, 2009.
6. P. Fogarty. *Catching the fire on grids*. PhD thesis, The University of Vermont, 2003.
7. D. B. Fogel. *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*. IEEE Press, Piscataway, NJ, USA, 1995.
8. A. Gilbers and R. Klein. A new upper bound for the vc-dimension of visibility regions. *Comput. Geom.*, 47(1):61–74, 2014.
9. J. H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. University of Michigan Press, Ann Arbor, MI, 1975.
10. R. Klein and E. Langetepe. Computational Geometry Column 63. *SIGACT News*, 47(2):34–39, June 2016.
11. K. Ng and P. Raff. Fractional firefighting in the two dimensional grid. Technical report, DIMACS Technical Report 2005-23, 2005.
12. I. Rechenberg. *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Number 15 in *Problemata*. Frommann-Holzboog, Stuttgart-Bad Cannstatt, 1973.
13. H.-P. Schwefel. *Numerical Optimization of Computer Models*. John Wiley & Sons, Inc., New York, NY, USA, 1981.
14. P. Wang and S. A. Moeller. Fire control on graphs. *Journal of Combinatorial Mathematics and Combinatorial Computing*, 41:19–34, 2002.