

Complexity

The complexity of the Dyson Telescopes puzzle

ERIK D. DEMAINE, MARTIN L. DEMAINE, RUDOLF FLEISCHER,
ROBERT A. HEARN, AND TIMO VON OERTZEN

ABSTRACT. We give a PSPACE-completeness reduction from QBF (quantified Boolean formulas) to the Dyson Telescopes puzzle where opposing telescopes can overlap in at least two spaces. The reduction does not use tail ends of telescopes or initially partially extended telescopes. If two opposing telescopes can overlap in at most one space, we can solve the puzzle in polynomial time by a reduction to graph reachability.

1. Introduction

The complexity of many motion-planning problems has been studied extensively in the literature. This work has recently focused on very simple combinatorial puzzles (one-player games) that nonetheless exhibit the theoretical difficulty of general motion planning; see, e.g., [1]. Two main examples of this pursuit are a suite of pushing-block puzzles, culminating in [2; 3], and a suite of problems involving sliding-block puzzles [4]. In pushing-block puzzles, an agent must navigate an environment and push blocks in order to reach a goal configuration, while avoiding collisions. The variations of pushing blocks began with several versions that appeared in video games (the most classic being Sokoban), and continued to consider simpler and simpler puzzles with the goal of finding a polynomially solvable puzzle. Nonetheless, all reasonable pushing-block puzzles turned out to be NP-hard, and many turned out to be PSPACE-complete, with no problems known to be in *NP*, except in one trivial case where solution paths are forced to be short. Similarly, sliding-block puzzles are usually PSPACE-complete, even in very simple models.

Fleischer's work was partially supported by a grant from the National Natural Science Fund China (grant no. 60573025).

In this paper we consider a motion-planning puzzle, the Dyson Telescopes puzzle. It takes the form of an enjoyable computer game [5], invented and developed by the Dyson company to advertise a vacuum cleaner called “Telescope” that is retractable like an astronomical telescope. The puzzle is perhaps most closely related to sliding blocks, in the sense that the agent is outside the environment. At any time, the agent can extend or retract one of several “telescopes”, each of which has a specified, fixed length in extended form. Erickson [6] posed the complexity of the problem in 2003. The complexity remained open despite fairly extensive pursuit—it seemed nearly impossible to build gadgets that required multiple entrances. Thus we hoped that it would be the first “interesting” yet polynomially solvable motion-planning puzzle.

We prove that the Dyson Telescopes puzzle is indeed polynomially solvable in a fairly natural situation in which the extended forms of opposing telescopes (two telescopes on the same row or column, pointing towards each other) overlap in at most one space. However, some of Dyson’s puzzles do not satisfy this restriction. We prove that this small flexibility in the general form of the problem in fact makes the problem PSPACE-complete.

The polynomial-time algorithm for the restricted form of the telescopes game is particularly interesting because such puzzles are nonetheless enjoyable for humans to play. All but a few of the hundreds of levels of the puzzle on the Dyson homepage [5] (mainly the Grandmaster levels) do not have opposing telescopes that overlap in more than one square. Therefore we expect that our algorithm can be used to design enjoyable instances of the telescope game, enumerating over puzzles within this restricted family (either by hand or by some automatic process), and automatically computing which puzzles are solvable. Our algorithm can also find the *shortest* solution, for most reasonable weighting functions, enabling the puzzle designer to find the *hardest* puzzle according to a particular difficulty measure, such as the solution requiring the longest sequence of moves or requiring a “difficult to see” sequence of moves.

1.1. Description of the problem. In the Dyson Telescopes puzzle, the goal is to maneuver a ball on a two-dimensional square grid from a starting position to a goal position, by extending and retracting telescopes on the grid; refer to Figure 1. An instance of the problem consists of an $n \times m$ grid, a number of telescopes on this grid, and the ball’s starting position and goal position. Each telescope is specified by its position, its direction (up, right, down, left), and its length, i.e., the number of spaces it can be extended. Each telescope can be in either an *extended* or a *retracted* state. Initially, all telescopes are retracted. A move is made by changing the state of a telescope.

If a telescope is extended, it will expand in its direction until it is blocked (i.e., there is a telescope occupying the space where the telescope would extend

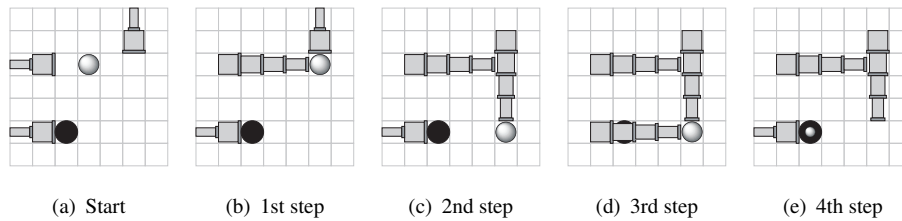


Figure 1. This example depicts a sample situation from the original game where all telescopes have length 3. We can solve this instance as follows: We extract the first telescope to push the ball to the right, where we then can push it downwards into the row of the lower telescope; when we extend and retract the lower telescope, it will finally pull the ball back to the goal position.

to next), or until it reaches its full length. If a ball blocks the extension of the telescope, the ball is pushed in the direction of the telescope, either until it is blocked by another telescope or until the pushing telescope is fully extended; see Figure 1(d). On the back side of the telescope (i.e., in the opposite direction as the telescope extends), there is a one-space tail. When the telescope is extended, this tail is retracted.

If an extended telescope is retracted, it is retracted all the way until it occupies only its base space. If the space behind the telescope is not occupied, its tail will be extended and occupy this space (and possibly push the ball). If the telescope end touches the ball when being retracted, it pulls the ball with it, so that the ball will move to the position directly in front of the retracted telescope; see again Figure 1(d).

We prove that it is PSPACE-complete to determine whether a given problem instance has a series of telescope movements that moves the ball from the starting position to the goal position (think of the goal square as a hole; the ball will fall down as soon as it is pushed across the goal square). We do this by constructing a circuit solving QBF, using gadgets of telescope configurations to simulate Boolean variables, logical gates, etc. If opposing telescopes are not allowed to overlap in more than one space, we give a polynomial time algorithm to find a solution.

Alternative versions of the game allow the telescopes to be partially extended in the initial state, or to not consider a tail end of the telescopes. We show that these modifications do not change the complexity of the problem.

2. Gadgets used in the reduction

In this section, we introduce various gadgets made from configurations of telescopes. These gadgets usually have some entrances and exits labeled by capital letters. We usually describe all the possible paths along which the ball can travel from an entrance to an exit.

2.1. Basic gadgets. We use the symbols in Figure 2 for simple tracks, simple crossings, division of the path, and union of paths, which are easy to implement. We assume that passage through one-way, split and join gadgets is possible only in the appropriate directions. Figure 3 shows the join and split gadgets.




Figure 2. Simple gadgets.

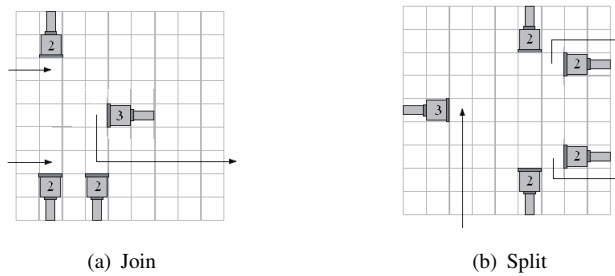


Figure 3. Details of the join and split gadgets.

Figure 4 shows a pair of *opposing telescopes* (the number on a telescope indicates its length). The pair is said to be *active* if one of the telescopes is extended with its end between the black and the white dot, and *inactive* otherwise.

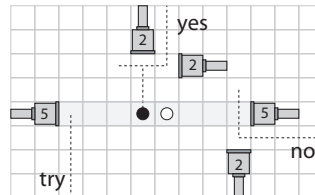


Figure 4. Opposing telescopes.

If the pair is inactive and the ball enters from *try*, it can only leave the gadget at *no*. On its way from *try* to *no*, it may activate the pair as follows. First, we retract all telescopes in the gadget. Then we extend the left telescope to full length (so that it covers the white dot square) and extend the right telescope until it is blocked by the left telescope just to the right of the white dot square (i.e., we extend it by three spaces). Then we retract the left telescope, put the ball into the gadget along the *try* path, and push it to the white dot square. Then we pull the ball to the *no* exit by retracting the right telescope. Note that this action leaves the pair in an active state.

If the pair is active and the ball enters from *try*, then the right telescope must be extended to just cover the white dot square. Then we can push the ball to the black dot square, where it can be picked up by the top telescopes so that it can leave the gadget at *yes*. We can also leave the pair at the *no* exit. In both cases, we may choose to leave the opposing pair either active or inactive.

Note that the ball can exit the gadget via *yes* and *no*, but it cannot enter the gadget at these points. We may lengthen the left and right telescopes (increasing the size of the gaps) and vertically flip the sides of the *try*, *yes*, or *no* pathways without changing the properties of the gadget, as long as we maintain the two-space overlap of the left and the right telescopes.

2.2. Variable gadget. When we move the ball from *try* to *yes* in an active opposing pair we may leave the pair active. To force it become inactive, we construct a *reset gadget*, shown schematically in Figure 5. Each grey rectangle represents an opposing pair. A single telescope extends along the lower pathway *r*, crossing the path of the lower telescope of β and ending in the path of the lower telescope of γ . The ball cannot enter β directly; it must first enter γ .

LEMMA 1. *The ball can always move through a reset gadget from in to out, but this forces the opposing pair α to become inactive.*

PROOF. The ball can only pass along path *r* if the lower telescopes of both β and γ are retracted. If an opposing telescope is also retracted, the corresponding pair will become inactive. If both upper telescopes of β and γ are extended to keep the opposing pairs active, α must be inactive. Since the ball can leave the gadget only if both β and γ are active, this is only possible if α is inactive. Note that the initial states of the opposing pairs are not important because we can activate β and γ (deactivating α). \square

We attach three independent reset gadgets to a single opposing pair α to construct a *variable gadget*, shown in Figure 6. Here, each pair (β_i, γ_i) corresponds to one reset gadget; the internal reset pathways are not shown. This is our workhorse gadget, forming the basis of all the following constructions. We say that the variable gadget is *open* (*closed*) if α is active (inactive).

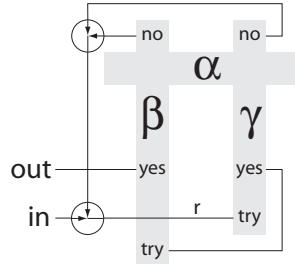


Figure 5. Reset gadget.

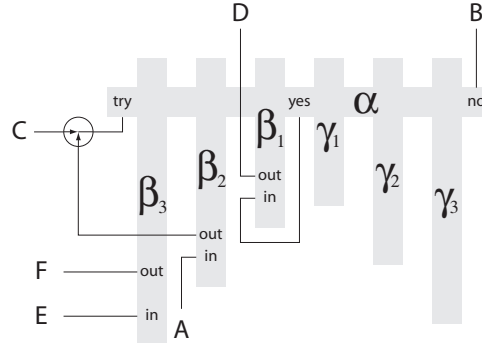


Figure 6. Variable gadget.

LEMMA 2. *In a variable gadget, traversal from either A or C to B is always possible, and may open the gadget. Traversal from C to D is possible precisely if the gadget is open, and forces it to close. Traversal from E to F is always possible, and forces the gadget to close. No other traversals are possible.*

PROOF. By properties of opposing telescopes, Lemma 1, and the pathways shown in Figure 6. \square

2.3. 3SAT gadget. Given a 3CNF formula W (a propositional formula in conjunctive normal form with three disjuncts per clause) with m clauses and n variables, we construct a *3SAT gadget*, shown in Figure 7, to test the formula. We use an $m \times 3$ array of variable gadgets. The three gadgets in row i correspond to the variables in clause i .

For each variable v and truth value $b \in \{0, 1\}$, we connect the A - B lines of all variable gadgets corresponding to $v = b$ into a chain. We also connect the E - F lines of all variable gadgets corresponding to $v = 1 - b$ into another chain. We concatenate these two chains by joining the last B line of the first chain to the first E line of the second chain. Finally we connect the first A line of the chain to an input channel $(v = b)_{in}$, and the last F line in the chain to an output channel $(v = b)_{out}$ of our 3SAT gadget.

We connect together the D lines of the three variable gadgets on row i and the C lines of the three variable gadgets on row $i + 1$, so that it is possible to go from any of the three D lines to any of the three C lines. We connect an input channel $test$ to the C lines of row 1. We connect the D lines of row m to an output channel $pass$.

Thus, the 3SAT gadget has $4n + 2$ ports ($in(v = b)$ for each v and b , one test input, and as many outputs).

LEMMA 3. *Consider a 3SAT gadget for a formula W . If the ball enters at $(v = b)_{in}$, it can only exit the gadget at $(v = b)_{out}$. This may open all gadgets*

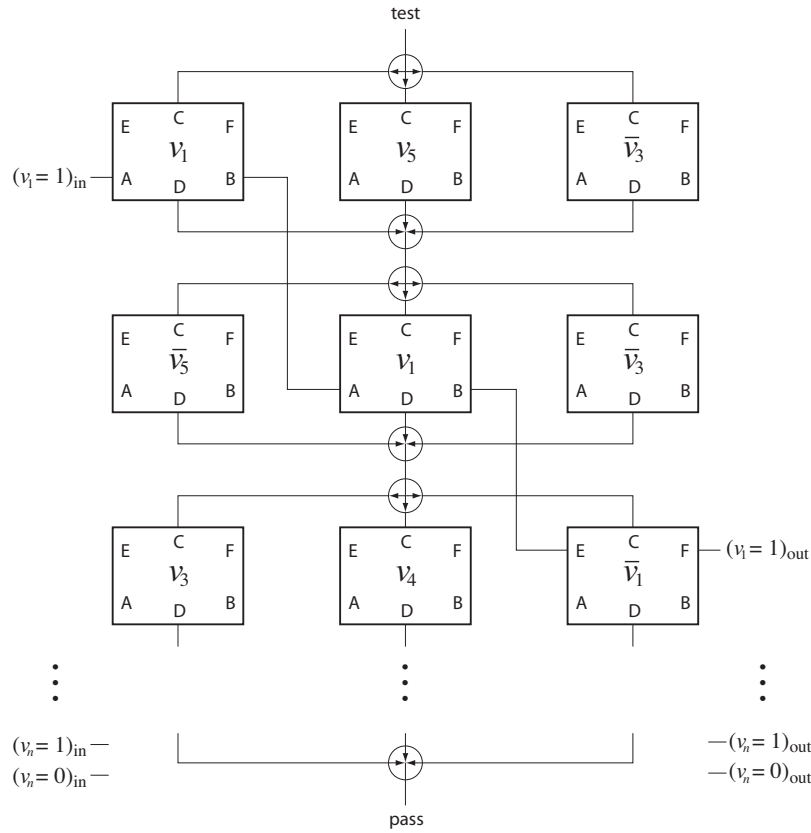


Figure 7. A 3SAT gadget. Shown are the test path and the path $(v_1 = 1)_{in}$ to $(v_1 = 1)_{out}$, where v_1 appears only in the first three clauses (twice positive, once negated).

corresponding to $v = b$ and must close all gadgets corresponding to $v = 1 - b$. The ball can also move from test to any $(v = b)_{out}$, and this must close all gadgets corresponding to $v = 1 - b$.

There exists an assignment $v_1 = b_1, \dots, v_n = b_n$ satisfying W if and only if the ball can traverse the gadget from test to pass (after first traversing it from $(v_i = b_i)_{in}$ to $(v_i = b_i)_{out}$, for $i = 1, \dots, n$).

PROOF. If the ball enters at $(v = b)_{in}$, it first reaches a chain of A - B channels through variable gadgets. It must follow the chain because in a variable gadget the only way from A leads to B . This may open all these gadgets. After the chain of A - B channels, the ball must traverse a chain of E - F channels which is also possible in only one way. This forces the corresponding variable channels to close.

If the ball enters at *test*, it follows a chain of *C-D* channels through variable gadgets. It may exit a variable gadget corresponding to $v = b$ at *B* and then follow the chain of *A-B* channels as above. It may open some gadgets corresponding to $v = b$, but then must close all gadgets in the *E-F* chain corresponding to $v = 1 - b$. This ensures that no variable is assigned more than one truth value (i.e., if any variable gadget corresponding to $v = b$ is open, then all variable gadgets corresponding to $v = 1 - b$ are closed, and vice versa).

So, if a path from *test* to *pass* of open variable gadgets exists, the corresponding variable assignment satisfying W can be read off. On the other hand, for each solution $v_1 = b_1, \dots, v_n = b_n$ of W , the n traversals from $(v_i = b_i)_{in}$ to $(v_i = b_i)_{out}$ are possible, opening a path from *test* to *pass*. \square

3. PSPACE-completeness

We now show that the Dyson Telescopes puzzle is PSPACE-complete. It is easy to see that the problem is in *PSPACE*, since the state of all telescopes and the ball position can be stored in linear memory. To show hardness, we reduce the problem from Quantified Boolean Formulas (QBF).

3.1. Countdown unit. We need a *countdown unit* that can be traversed at most 2^n times. The gadget is shown in Figure 8. We chain together $n + 1$ variable gadgets, linking each gadget's *B* exit to the next gadget's *C* entry. We combine the *D* exits into an overall exit line, and link the last variable gadget's *B* exit to another exit of the countdown unit.

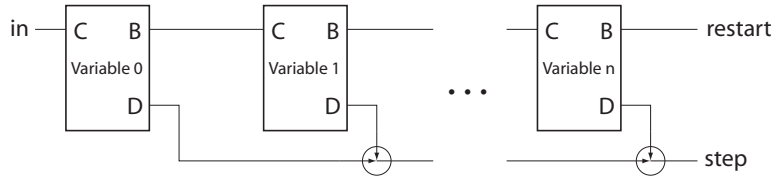


Figure 8. Countdown unit.

LEMMA 4. *When the ball enters the countdown unit for the first time, it can leave it at restart. After the gadget has been traversed from in to restart, it can be at most 2^n times traversed from in to step, before it must again be traversed from in to restart.*

PROOF. If all variable gadgets are closed, the ball can only leave them at *B*. After moving from *in* to *restart*, all or some of the gadgets may be open. But then the *in-step* channel can be used at most 2^n times, as can be seen by induction. \square

3.2. Reduction from QBF. Let

$$W = \forall v_{1,1} \exists v_{1,2} \forall v_{2,1} \exists v_{2,2} \dots \forall v_{n,1} \exists v_{n,2} f(v_{1,1}, \dots, v_{n,2})$$

be a quantified boolean formula with a 3CNF formula f . We build a gadget to test W using a 3SAT gadget for f , one countdown unit of size n , and a chain of n additional variable gadgets. The construction is shown in Figure 9.

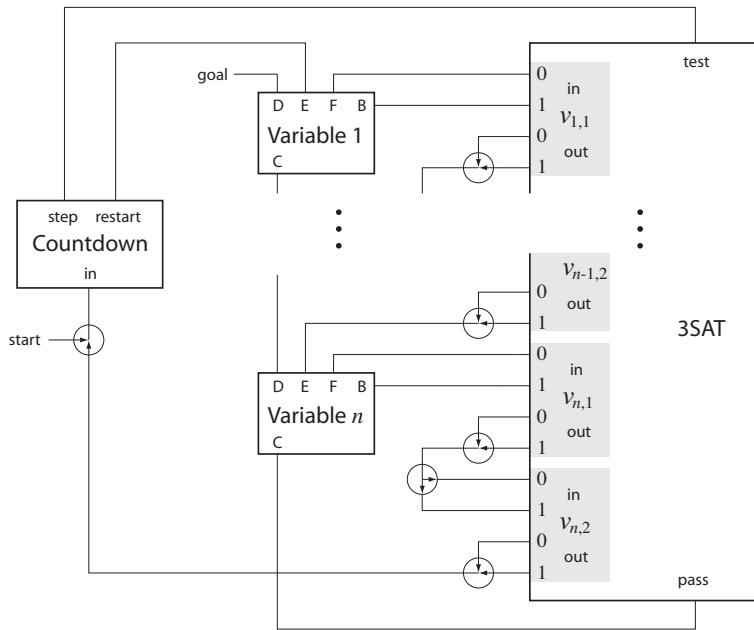


Figure 9. Reduction from QBF.

Each D exit of the variable gadgets is linked to the C entry of the previous gadget. Each F exit is linked to the E entry of the next variable gadget, however not directly but via (1) the $(v_{i,1} = 0)$ channel of the 3SAT gadget, and then (2) either the $(v_{i,2} = 0)$ or the $(v_{i,2} = 1)$ channel of the 3SAT gadget. The B exit of each variable gadget is also linked to the E entry of the next gadget, via (1) the $(v_{i,1} = 1)$ channel of the 3SAT gadget, and then (2) either the $(v_{i,2} = 0)$ or the $(v_{i,2} = 1)$ channel of the 3SAT gadget. The $(v_{n,2} = 0)$ and $(v_{n,2} = 1)$ channels of the 3SAT gadget are linked to the in entry of the countdown unit, whose $step$ exit is linked via the $test$ channel of the 3SAT gadget to the last variable gadget's C entry. The first gadget's D exit is linked to the $goal$. The starting point is also linked to the in entry of the countdown unit. The $restart$ exit of the countdown unit is linked to the first variable gadget's E entry point.

THEOREM 5. W is true if and only if the ball can move from $start$ to $goal$.

PROOF. We first describe how we can systematically test the formula for all possible truth assignments according to the quantifiers in W .

Initially, we must traverse the countdown unit from *in* to *restart*. Whenever the ball leaves the countdown unit at *restart*, it institutes a restart of the variable gadgets: all of them must be passed from E to F , so all of them are closed, and all variables with universal quantifiers are set to 0; all other variables can be chosen freely.

Next we can test the 3SAT gadget with the current choice of variables truth assignment. If we can pass the gadget successfully, the ball ends up at the C entrance of the gadget of the last variable n . Since this gadget is closed, the ball can only leave it at B , and we open the gadget while passing through. Since we leave the gadget at B we can now set $v_{n,1}$ to 1 and then choose a new arbitrary value for $v_{n,2}$. Then we test the 3SAT gadget again with this new truth assignment. But this time we can leave the gadget for variable n at D , pass through the gadget for variable $n - 1$, opening it, and set $v_{n-1,1}$ to 1. Then we can choose a new value for $v_{n-1,2}$, traverse the gadget for variable n along $E-F$, thereby closing it, and reset $v_{n,1}$ to 0. Finally we can choose a new value for $v_{n,2}$.

In this way, the chain of variable gadgets enumerates all possible settings of variables with universal quantifiers. Whenever we open a variable gadget, its corresponding \forall -variable is set to 1, and whenever we use the $E-F$ channel to close the gadget, we reset its \forall -variable to 0. For the corresponding \exists -variables we can choose arbitrary values. A gadget can only be opened if all gadgets below (i.e., with higher index) have already been opened, so the gadgets act as a counter which must be passed at least 2^n times to reach the goal.

Every time this counter is increased (i.e., reaches the entry of the countdown unit), it must pass the countdown unit and the 3SAT test channel. If the ball were to traverse the 3SAT unit from *test* to any $(v = b)_{out}$, then the countdown unit would have to be passed more often than the counter given by the additional variable gadgets. But these must be passed 2^n times to reach *goal*. Since the countdown unit does not allow more than 2^n traversals from A to B , it would have to be left at *restart* before we reach *goal*, which would reset the whole structure. Therefore, the ball cannot move from *test* to any other exit than *pass* if it wants to reach *goal*.

By the same argument, whenever a variable gadget is traversed from C to B , it must be opened, otherwise more than 2^n passages are required to reach *goal*, and the whole structure is reset.

If the 3SAT gadget is tested with every possible variable setting for the variables with universal quantifiers and a choice of values for the variables with existential quantifiers, W is true. If on the other hand W is true, there is such a

selection for each possible setting for variables with universal quantifiers, and a path from *start* to *goal* exists. \square

4. Opposing telescopes that overlap in at most one space

In this section, we show that the Dyson Telescopes puzzle is in P if opposing telescopes cannot overlap in more than one space. Let D denote an instance of such a problem, and let T_1, \dots, T_n denote the telescopes. A *constellation* of the telescopes is an assignment of integers to the telescopes describing how far the telescopes are extended.

A *direct traversal* from T_i to T_j is a sequence of telescope extensions and retractions such that the ball is initially attached to T_i , finally attached to T_j , and in between it is not pushed or sucked by any other telescope. A *traversal* from T_1 to T_n is a sequence of direct traversals, where the ball is first attached to T_1 and ends up attached to T_n .

We first assume that D has no opposing pairs. Then we can define an induced directed graph G_D with the telescopes as vertices and an edge from T_i to T_j if

- T_i and T_j are orthogonal. Let f be the space in which they overlap.
- T_i and T_j can be extended at least up to the space before f .
- f is either the first space in front of T_i , the first space not reachable by T_i (i.e., the space to which the ball would be pushed if T_i was completely extended), or there is a third telescope T_k that can be extended to the space after f in the extension path of T_i .

LEMMA 6. *Assume D has no opposing pairs. If the ball is attached to a telescope T_i , then a direct traversal from T_i to another telescope T_j is possible precisely if there is an edge from T_i to T_j in G_D , independent of the current constellation of D .*

PROOF. If there is an edge from T_i to T_j , then obviously a direct traversal is possible.

Assume a direct traversal is possible in some constellation. Then T_i and T_j must be able to reach a common space f . Since there are no opposing pairs of telescopes, T_i and T_j must be orthogonal. Then, f is the only space reachable by both telescopes. We can transfer the ball from T_i to f if f is either the first or last reachable space of T_i , or if the space after f in the path of T_i is blocked by another telescope T_k . In any case, the edge (T_i, T_j) exists in G_D . \square

Now assume D contains an opposing pair as shown in Figure 10, where A and B overlap in at most one space, denoted by the black dot (if it exists). There may also be a third telescope C pointing to this space (or even extending beyond). There might even be a fourth telescope (not shown, it can be handled

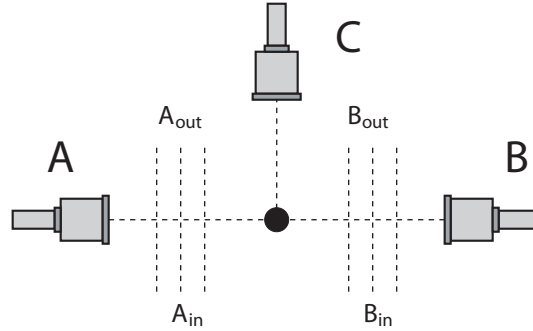


Figure 10. Opposing telescopes with at most one overlapping space.

analogously) opposing C and extending up to or beyond the black dot space. We define the graph G_D as before, but for the opposing pair we must add some additional edges as described below. For a telescope T , let T_{in} denote the set of all telescopes with an edge pointing to T , and T_{out} the set of telescopes to which T points in G_D . Note that C may or may not be in A_{in} and B_{in} , depending on the overall configuration of the at most four telescopes covering the black dot square and the initial position of the ball. Actually, for the construction below it is sufficient to assume that C is not in A_{in} and B_{in} .

LEMMA 7. *Traversal from any telescope in $A_{in} \cup B_{in}$ to C (if it exists) and any telescope in $A_{out} \cup B_{out}$ is possible in every constellation of D .*

PROOF. If C exists, we can move the ball from any telescope $T \in A_{in}$ to C as follows. First, we retract A , B , T , and C . Then we extend A completely. If we now extend B , it will be stopped just right of the black dot. Now we can retract A and move the ball from T to the line of A , which is possible since there is an edge from T to A in G_D . If we then extend A , the ball will come to rest on the black dot, where we can pick it up with C .

All other traversals are trivially possible. \square

Although the traversal from T to C in the proof above is done via A , it is impossible to traverse directly from A to C without prior preparation of the opposing pair if C extends exactly to the square above the black dot square. If the ball is initially placed in the opposing pair and the pair is not initially set up such that traversal to C is possible, C cannot be reached directly. This means we should add the following edges to G_D for each opposing pair (A, B) with one space overlap (and maybe an orthogonal telescope C pointing to the overlap space):

- edges $A \rightarrow B$ and $B \rightarrow A$;
- edges $T \rightarrow C$ for all $T \in A_{in} \cup B_{in}$, if C exists;

- edge $A \rightarrow C$, if C exists and can be extended to block A or B , or B is initially extended immediately to the right of the overlap space;
- edge $B \rightarrow C$, if C exists and can be extended to block A or B , or A is initially extended immediately to the left of the overlap space.

Note that in the second case the edges $T \rightarrow C$ are a shortcut for $T \rightarrow A \rightarrow C$ because we do not always want to add edge $A \rightarrow C$ to the graph (depending on the initial placement of the ball).

LEMMA 8. *Let D be an instance of the Dyson Telescopes puzzle with no opposing pair having more than one space overlap. Let G_D be the induced graph with edges as described above. Then, D has a solution exactly if there exists a path in G_D from a telescope that reaches the starting position of the ball to a telescope that reaches the goal position.*

PROOF. If there is a sequence of telescope movements that move the ball from start to goal, this induces a sequence of telescopes. If the start position of the ball is within an opposing pair (A, B) and both telescopes are initially retracted, the ball cannot leave the segment between A and B via C . But paths from A and B to all nodes of A_{out} and B_{out} exist in G_D , so the first telescope moves until the ball leaves the segment between A and B are reflected by edges in G_D . If the ball starts within the pair and one of the telescopes is not extended such that leaving at C would be possible, this is also reflected in G_D . Afterwards, all direct traversals of the winning strategy correspond to edges in G_D .

If on the other hand a path in G_D exists, it can easily be translated to a sequence of ball traversals (either direct or through opposing pairs) that gives a strategy to move the ball from start to goal. \square

COROLLARY 9. *The Dyson Telescopes puzzle is in P if opposing telescopes can overlap in at most one space.*

5. Summary and outlook

We showed that, in general, the problem of deciding whether the ball can move from start to goal in a setting of the Dyson Telescopes puzzle is PSPACE-complete. We also gave a polynomial-time algorithm if opposing pairs are restricted to at most one space of overlap.

Both the PSPACE-completeness proof and the algorithm for the restricted case also work if the back ends of the telescopes are taken into account and if the telescopes can initially be arbitrarily (partially) extended. Note that the PSPACE-hardness proof requires rather long telescopes. It would be interesting to investigate the complexity status of the problem with bounded-length telescopes.

Acknowledgements

We thank Jeff Erickson for posing the problem to us, and Eva M. Jungen for helpful discussions on the polynomial-time algorithm. We also thank the anonymous reviewer for his helpful suggestions.

References

- [1] Erik D. Demaine. Playing games with algorithms: Algorithmic combinatorial game theory. In *Proc. 26th MFCS 2001*, pp. 18–32, LNCS 2136, 2001.
- [2] Erik D. Demaine, Martin L. Demaine, Michael Hoffmann, and Joseph O’Rourke. Pushing blocks is hard. *Computational Geometry: Theory and Appl.* 26(1):21–36, 2003.
- [3] Erik D. Demaine, Robert A. Hearn, Michael Hoffmann. Push-2-F is PSPACE-complete. In *Proc. of the 14th Canadian Conference on Computational Geometry (CCCG’02)*, pp. 31–35, 2002. <http://www.cs.uleth.ca/~wismath/cccg/papers/31.ps>.
- [4] Robert A. Hearn, Erik D. Demaine. PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation. *Theoretical Computer Science*, to appear. <http://www.arXiv.org/abs/cs.CC/0205005>
- [5] Dyson UK’s Telescope game. To be found at <http://www.dyson.co.uk/game/>.
- [6] Jeff Erickson. Personal communication, September 2003.

ERIK D. DEMAINÉ
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY
CAMBRIDGE, MA
UNITED STATES
edemaine@mit.edu

MARTIN L. DEMAINÉ
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY
CAMBRIDGE, MA
UNITED STATES
mdemaine@mit.edu

RUDOLF FLEISCHER
FUDAN UNIVERSITY
SHANGHAI KEY LABORATORY OF INTELLIGENT INFORMATION PROCESSING
DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING
SHANGHAI
CHINA
rudolf@fudan.edu.cn

ROBERT A. HEARN
NEUKOM INSTITUTE FOR COMPUTATIONAL SCIENCE, DARTMOUTH COLLEGE
SUDIKOFF HALL, HB 6255
HANOVER, NH 03755
UNITED STATES
robert.a.hearn@dartmouth.edu

TIMO VON OERTZEN
MAX-PLANCK-INSTITUTE FOR HUMAN DEVELOPMENT
BERLIN
GERMANY
vonoertzen@mpib-berlin.mpg.de

