Pushing Blocks is NP-Complete for Noncrossing Solution Paths

Erik D. Demaine* Michael Hoffmann†

Abstract

We prove NP-hardness of a wide class of pushing-block puzzles like the classic Sokoban, generalizing several previous results [3, 4, 6, 7, 12, 14]. The puzzles consist of unit square blocks on an integer lattice; all blocks are movable. The robot may move horizontally and vertically in order to reach a specified goal position. In the PUSH-k puzzle, the robot can push up to k blocks in front of it as long as there is at least one free square ahead. Other variations were introduced to make puzzles more tractable, in which blocks must slide their maximal extent when pushed (PUSH-PUSH), and in which the robot's path must not cross itself (PUSH-X). We prove that all of these puzzles are NP-hard.

Keywords: Motion planning, combinatorial games, computational complexity.

1 Introduction

Algorithmic motion planning is a large area of computational geometry with applications in robotics, assembly planning, and computer animation; see, e.g., [13] for a survey. The standard type of problem involves moving a robot from one configuration to another while avoiding fixed obstacles. A recent direction introduced by Wilfong [14] is a class of problems in which robots are permitted to move some of the obstacles in order to increase maneuverability. As robots become more powerful at manipulation, an understanding of such models becomes increasingly important. Current-day applications include automated warehouse control and warehouse navigation; see, e.g., [8]. A representative abstraction of such applications is the popular Sokoban puzzle [2, 7], which is known to be PSPACE-complete [2]. In this paper we study several variations of simpler puzzles, and show all of these models are NP-hard using a reductions from 3-coloring of planar graphs [9]. Some variations are additionally known to be NP-complete, others PSPACE-complete, while the complexity of most variations is unresolved between NP and PSPACE.

Problems. Our hardness results are particularly surprising given the simplicity of the model of motion and obstacle manipulation. Consider a rectangular $n \times m$-grid in which each square is marked either *free* or *blocked*. A *robot* can move horizontally and vertically in the grid, and thereby push up to k blocks in front of it, for some constant k. See Figure 1, in which the blocked positions are shaded and the robot is shown as circle, pushing two blocks.

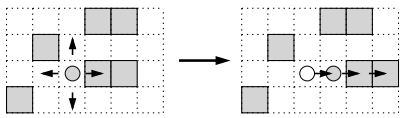

Figure 1: Example of pushing blocks.

The PUSH-k problem [3, 4] is to decide whether there is a sequence of moves starting at a specified free position and ending at a specified goal position. If we omit the restriction on how many blocks the robot can push at once (i.e., $k = \infty$), we obtain the problem PUSH-* [1, 5, 6, 11].

To make these basic problems more computationally tractable, we can impose additional simplifying constraints on the robot's motion. The PUSH-PUSH model [3, 4, 12] requires that, once a block is pushed, it slides the maximal extent in that direction. This model can be thought of representing either sliding blocks on a frictionless surface, or the situation in which blocks cannot be pushed by precise amounts but can be consistently pushed against other blocks. The PUSH-X model [3] requires that the robot not cross its own path. This model is not intended to represent reality, but rather was proposed to severely restrict the motions in order to make the problem computationally tractable. In particular, any version of PUSH-X is automatically in NP, unlike the general problem.

Related Work. Out of these many problem variations, several individual cases have been studied. The original paper by Wilfong [14] studies a more flexible model in which the blocks can be more general than squares, and the robot can both push and pull blocks. Dhagat and O'Rourke [6] initiated the PUSH- line of models, and proved that PUSH-* is NP-hard if some blocks can be tied to the board, making them unpushable. This result was later strengthened to PSPACE-completeness [1].

Our models disallow blocks from being tied to the board, making the problem easier to solve and hardness gadgets more difficult to construct. Progress in this setting has only been made in the last two years, during which NP-hardness was proved for PUSH-PUSH-1 [3, 4, 12], PUSH-1 [3], and PUSH-* [11]. These results leave two aspects open: PUSH-k for general k, and the PUSH-X noncrossing restriction. None of the pre-

*Department of Computer Science, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada, eddemaine@uwaterloo.ca.

†Institute for Theoretical Computer Science, ETH Zurich, CH-8092 Zurich, Switzerland, hoffmann@inf.ethz.ch.

vious reductions extend easily to these scenarios, except for the PUSH-* reduction [11] which also applies to PUSH-*-X.

Results. In this paper, we provide one construction that settles the NP-hardness of PUSH-k, PUSH-k-X, PUSH-PUSH-k, and PUSH-PUSH-k-X, for any fixed $k \in \mathbb{N}$. In particular, our results subsume a number of previous NP-hardness proofs of pushing-block variations [3, 4, 6, 7, 12, 14], and solves the PUSH-X open problem from [3]. Thus, this paper together with [11] prove that all pushing-block puzzles described above are NP-hard. The new idea in our reduction is to force the robot to follow constrained Eulerian tours of planar graphs and carry a constant amount of information along each edge of the graph. This idea is in contrast to all previous approaches of building circuits based on graphs, which seem to inherently require crossings.

2 The Reduction

We first consider the problem PUSH-1-X, in which the robot is restricted to push only one block at any time. Later we will describe how to modify the construction for $k > 1$.

Our reduction is from planar 3-coloring, which is known to be NP-complete, even if no vertex has degree larger than four [9]. We are given a planar embedding of a connected undirected graph $G = (V, E)$, and will construct a PUSH-1 puzzle from G that is solvable if and only if G is (vertex) 3-colorable.

Let $\overrightarrow{G}$ be the directed planar graph resulting from G by replacing each undirected edge by two directed edges of opposite orientation. By a well-known theorem (cf. [10]), there is a Eulerian tour in $\overrightarrow{G}$, since for every vertex the number of incoming edges equals the number of outgoing edges. Moreover, we claim that there is always a *planar* Eulerian tour T in $\overrightarrow{G}$, where planar means that it does not even cross itself at vertices, i.e., can be drawn with a pencil in one piece without crossing. (The fact that edges do not intersect in their interior is already implied by the planarity of G.)

Lemma 1 *There is a planar Eulerian tour T in $\overrightarrow{G} = (V, \overrightarrow{E})$.*

We omit the proof in this abstract.
The idea is to use T for traversing $\overrightarrow{G}$, thereby

- choosing the color of a vertex whenever leaving it and
- checking that the adjacent vertices are colored differently whenever traversing the second of the both edges from $\overrightarrow{G}$ representing an edge from G.

Of course, we must ensure that the colors chosen for the vertices are consistent, i.e., anytime we leave a specific vertex, the same color has to be chosen.

For this, we let two (directed) edges leaving the same vertex meet in what we call a *consistency gadget*, in order to assure that the same color is chosen in both. Of course, it might not be possible to join any pair of edges leaving a specific vertex this way without crossings, but it is sufficient to join the edges such that they form a tree under those junctions, and then all edges have to choose the same color by transitivity. Refer to the example in Figure 2: the edges of the graph G are drawn as dashed lines, the solid curve is a planar Eulerian tour of $\overrightarrow{G}$, and the wiggled segments denote the consistency junctions.

Similarly, any pair of edges representing the same edge of the original graph G are joined; these *coloring junctions* are shown as solid arrows in Figure 2. It is clear how to draw them, since both edges are always adjacent in the drawing. Also, coloring and consistency junctions do not interfere, since all of the latter are close to a vertex. By placing an appropriate gadget on the coloring junctions, we can forbid to choose the same color for adjacent vertices.

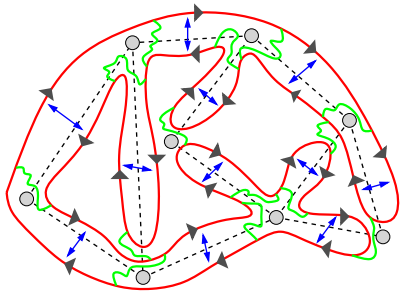

Figure 2: Example graph, tour and junctions.

The construction uses several gadgets, as described in Table 1. In the left column, we give a symbol to denote the type of gadget; in the right column we list the gadget's name, in parentheses its number of entries and exits, and a short description of its functionality. These gadgets are explained in detail in Sections 2.3–2.6.

	One-way gadget (1/1); can only be passed in direction of the arrow. However, once this happened, it is open, i.e. can be traversed in both directions arbitrarily.
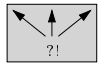	Fork gadget (1/3); can be left through any of the three exits, but only one of these paths may be entered.
	XOR-crossing gadget (2/2); a safe, leakage free crossing, provided that only one of the two paths is actually traversed.
	NAND gadget (2/2); joins two paths such that at most one of them can be traversed.

Table 1: The gadgets.

Since the robot can push at most one block, any two-times-two square of blocks is essentially fixed. Hence,

one can build the Eulerian tour as a PUSH-1 puzzle, using corridors that are surrounded by walls of thickness at least two. More precisely, each edge $\overrightarrow{e} = (u, v)$ from $\overrightarrow{G}$ will be represented by three corridors in the puzzle, as shown in Figure 3; the robot can choose to traverse any of them (but only one) through a fork gadget, this way assigning a color to the vertex u it leaves. The corridors of $\overrightarrow{e}$ are joined to the corridors of another edge f also leaving u (if existent) through a consistency gadget which guarantees that only the corridors of matching color can be traversed. Similarly, the corridors are joined to those corresponding to the opposite edge $\overleftarrow{e} := (v, u)$ through a coloring gadget, making sure that u and v are colored differently. Finally, the three paths rejoin, protected by one-way gadgets preventing the robot from stumbling backwards in the wrong direction.

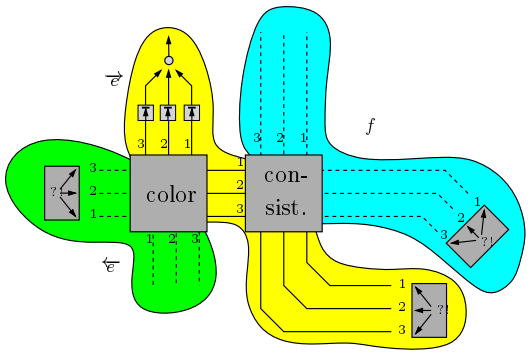

Figure 3: Construction for an edge.

2.1 The Coloring Gadget. This gadget is used to make sure that adjacent vertices do not get the same color. It joins two paths of three labeled corridors each in a symmetric way such that corridors with the same label are joined into a NAND gadget; see Figure 4.

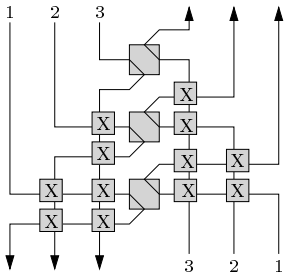

Figure 4: Coloring gadget.

Lemma 2 *Both sides of a coloring gadget can be traversed if and only if the robot chooses differently labeled paths.*

2.2 The Consistency-Check Gadget. This gadget is complementary to the coloring gadget. It is used to make sure that the robot cannot choose different colors for the same vertex, that is anytime it leaves the vertex, the same color has to be chosen. The gadget joins two paths of three labeled corridors each in a symmetric way such that corridors with the different labels are joined into a NAND gadget; see Figure 5.

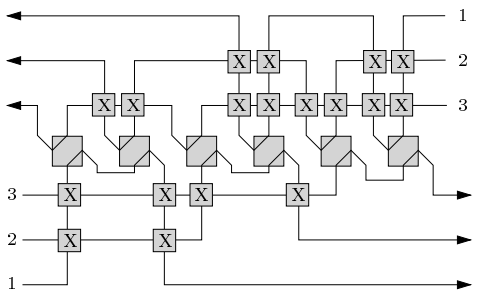

Figure 5: Consistency Gadget.

Lemma 3 *Both sides of a consistency gadget can be traversed if and only if the robot chooses paths with the same label.*

2.3 NAND Gadget. To make sure that only specific combinations of corridors are used by the robot, we need a gadget joining two corridors in such a way that at most one of them can be traversed: a NAND gadget. The layout is shown in Figure 6; it depends on whether both corridors are to be traversed in the same or in opposite directions.

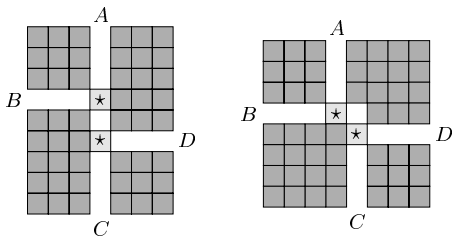

Figure 6: NAND gadgets.

Lemma 4 *Consider the gadgets shown in Figure 6.*

The left gadget can be traversed either $A \to B$ or $C \to D$ but not both ways.

The right gadget can be traversed either $B \to A$ or $C \to D$ but not both ways.

Remark 1 *For PUSH-k each marked block in the left gadget would have to be replaced by a sequence of $\lceil \frac{k+1}{2} \rceil$ blocks.*

Since it is always clear from the Eulerian tour in which direction an edge is to be traversed, there is no problem in choosing the appropriate type of NAND-gadget. Hence, we refer to this class of gadgets by the term "NAND-gadget" as if there was just a single one.

2.4 One-Way Gadget. The one-way shown in Figure 7(a) can be traversed in only one direction, from A to B. But note that once it has been traversed, it

is just an open corridor that can be traversed in both directions.

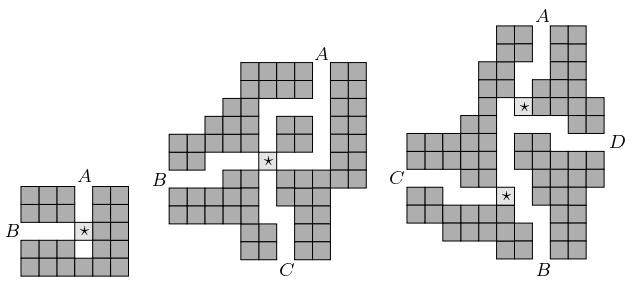

(a) One-Way (b) Fork (c) XOR-Crossing

Figure 7: The basic gadgets (after Figs. 2, 3, 10 of [3]).

2.5 Fork Gadget. The fork gadget shown in Figure 7(b) allows traversal from A to B or C, but once either of B or C has been reached, the other is inaccessible from A.

Note that it is possible to go e.g. from A to B and then from A to C, if the blocking square is pushed out of the corridor to C in between. But in our construction this can never occur, since the exits of the forked corridors are protected by one-way gadgets.

Also note that a three-way fork as in Table 1 is just a combination of two two-way forks.

2.6 XOR-Crossing Gadget. The XOR-crossing shown in Figure 7(c) can be traversed either from A to B or from C to D without leakage, i.e. it cannot be traversed from A to C or D or from C to A or B.

Note that the above statement is true for a single traversal only or more precisely, if the gadget is traversed only from either of A or C. This condition is fulfilled wherever XOR-crossings are used in our construction.

2.7 Main Theorems. In order to actually define the PUSH-1 puzzle, we need a start and a goal position for the robot. This is easily done by breaking the tour at some arbitrary vertex, yielding a path with start and goal position at its ends.

Theorem 5 PUSH-*1 is NP-hard.*

Theorem 6 PUSH-k *is NP-hard for any fixed* $k \in \mathbb{N}$.

Proof. As in Theorem 5. But corridors have to separated by walls of at least $k+1$ blocks and the NAND-gadget has to adopted as indicated in Remark 1. □

Corollary 7 PUSH-k-X, PUSH-PUSH-k, *and* PUSH-PUSH-k-X *are NP-hard.*

Proof. There is always a path for the robot that does not cross itself and wherever blocks are pushed, they are pushed as far as possible. □

3 Conclusion

We have shown NP-hardness of a broad class of pushing-block puzzles. Except for the noncrossing PUSH-X variants, it remains open whether these problem are in NP. Some of the problems might be PSPACE-complete, as has been shown for the fixed-block version of PUSH-* [1].

Another still open question is whether there is an "interesting" puzzle variant that is still tractable. Up to now, the only problem known to be in P is if the robot is restricted to monotonic paths [6], e.g., it can only move up and right.

Acknowledgments. We thank Martin Demaine and Joseph O'Rourke for helpful discussions.

References

[1] BREMNER, D., O'ROURKE, J., AND SHERMER, T. Motion planning amidst movable square blocks is PSPACE complete. Draft, 1994.

[2] CULBERSON, J. Sokoban is PSPACE-complete. In *Proc. Internat. Conf. Fun with Algorithms* (Elba, Italy, June 1998), Carleton Scientific, pp. 65–76.

[3] DEMAINE, E. D., DEMAINE, M. L., AND O'ROURKE, J. PushPush and Push-1 are NP-hard in 2D. In *Proc. 12th Canad. Conf. Comput. Geom.* (2000), pp. 211–219.

[4] DEMAINE, E. D., DEMAINE, M. L., AND O'ROURKE, J. PushPush is NP-hard in 2D. Tech. Rep. 065, Dept. Comput. Sci., Smith College, Northampton, Jan. 2000.

[5] DEMAINE, E. D., AND O'ROURKE, J. Open problems from CCCG'99. Tech. Rep. 066, Dept. Comput. Sci., Smith College, Northampton, Mar. 2000.

[6] DHAGAT, A., AND O'ROURKE, J. Motion planning amidst movable square blocks. In *Proc. 4th Canad. Conf. Comput. Geom.* (1992), pp. 188–191.

[7] DOR, D., AND ZWICK, U. SOKOBAN and other motion planning problems. *Computational Geometry: Theory and Applications 13*, 4 (1999), 215–228.

[8] EVERETT, H. R., AND GAGE, D. W. From laboratory to warehouse: Security robots meet the real world. *Internat. J. Robotics Research 18*, 7 (July 1999), 760–768.

[9] GAREY, M. R., AND JOHNSON, D. S. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman, New York, NY, 1979.

[10] HARARY, F. *Graph Theory.* Addison-Wesley, Reading, MA, 1972.

[11] HOFFMANN, M. Push-* is NP-hard. In *Proc. 12th Canad. Conf. Comput. Geom.* (2000), pp. 205–210.

[12] O'ROURKE, J., AND THE SMITH PROBLEM SOLVING GROUP. PushPush is NP-hard in 3D. Tech. Rep. 064, Dept. Comput. Sci., Smith College, Northampton, Nov. 1999.

[13] SHARIR, M. Algorithmic motion planning. In *Handbook of Discrete and Computational Geometry*, J. E. Goodman and J. O'Rourke, Eds., 1997, ch. 40, pp. 733–754.

[14] WILFONG, G. Motion planning in the presence of movable obstacles. *Ann. Math. Artif. Intell. 3* (1991), 131–150.