

PushPush- k is PSPACE-Complete[★]

Erik D. Demaine^a, Michael Hoffmann^b, and Markus Holzer^c

^a*Laboratory for Computer Science,
Massachusetts Institute of Technology,
200 Technology Square, Cambridge, MA 02139, USA
email: edemaine@mit.edu*

^b*Institute for Theoretical Computer Science,
ETH Zürich, CH-8092 Zürich, Switzerland
email: hoffmann@inf.ethz.ch*

^c*Institut für Informatik, Technische Universität München,
Arcisstraße 21, 80290 München, Germany
email: holzer@in.tum.de*

Abstract

We prove that a pushing-block puzzle called PUSH-PUSH- k is PSPACE-complete for any fixed $k \geq 1$. In this puzzle, a robot moves on a finite grid. Each grid position is either empty or occupied by a single obstacle block. While moving, the robot may push obstacle blocks in direction of its movement, subject to certain constraints. In particular, once an obstacle block starts moving, it continues to do so until it hits another obstacle or the grid boundary. The problem is to decide whether the robot can navigate from a given start position to a specified goal position on the grid.

1 Introduction

Pushing blocks. The generic pushing-block puzzle involves a rectangular grid of squares, each either empty, occupied by one of several *blocks*, or occupied by a unique *robot*. The robot can move horizontally or vertically to an empty square, or the robot can *push* an adjacent block or blocks in the direction of movement, subject to certain constraints. The goal of the puzzle is for the robot to reach a specified goal location, starting from a given configuration

[★] This work was initiated during Dagstuhl seminar No. 334 on “Algorithmic Combinatorial Game Theory,” held at Schloß Dagstuhl, Wadern, Germany, February 17–22, 2002.

of the board. The corresponding decision problem is to decide whether such a puzzle is solvable by a sequence of moves and pushes.

Puzzle Classification. In this paper, we consider the following two pushing constraints; see [Dem01] for a survey of other related constraints and results. The $\text{PUSH-}k$ constraint, for a fixed $k \geq 1$, is that the robot can push at most k consecutive blocks at once, and the PUSHPUSH constraint is that, when a block is pushed, it slides to its maximal extent, ending against another block or the boundary of the rectangular grid. Together, these two constraints are called $\text{PUSHPUSH-}k$. If the number of blocks which the robot can push at once is not restricted, we call the resulting constraint $\text{PUSHPUSH-}*$.

The $\text{PUSHPUSHPUSH-}k$ constraint is similar to $\text{PUSHPUSH-}k$, but blocks do not necessarily stop to move when another block is hit. Instead they continue to slide until the amount of blocks sliding simultaneously exceeds k or they hit the boundary.




Fig. 1. Example: a single move in different pushing models.

Related work. Several papers address specific pushing-block puzzles. It is easy to show that PUSHPUSH (and its variants) are in PSPACE . In particular, $\text{PUSH-}1$ was introduced by O’Rourke *et al.* [OS99], motivated by a computer game¹. They proved that the natural generalization of the $\text{PUSH-}1$ puzzle to three dimensions is NP-hard. The proof was extended to the two-dimensional version by Demaine, Demaine, and O’Rourke [DDO00b]. An alternate proof, which simultaneously resolved the less-constrained puzzle $\text{PUSH-}1$, was then published by the same authors [DDO00a]. At the same time, Hoffmann [H00] gave a construction showing that $\text{PUSH-}*$ is NP-hard. Subsequently, Demaine and Hoffmann [DH01] proved NP-hardness of $\text{PUSH-}k$ for any $k \geq 1$. In particular, they showed that $\text{PUSH-}1$ and $\text{PUSH-}1$ remain NP-hard even when the solution path must not cross itself. In this last case, called $\text{PUSH-}1\text{-X}$ and $\text{PUSH-}1\text{-X}$, the puzzles are also in NP, and hence are NP-complete. Without this noncrossing restriction, the puzzles do not obviously have polynomial-length solutions. Finally, Demaine, Hoffmann, and Hearn [DHH02] considered a model called $\text{PUSH-}k\text{-F}$ where some blocks

¹ See <http://www.pushpush.net> for a history of the game and its implementations.

may be marked as fixed, that is, immovable for the robot, and proved $\text{PUSH-}k\text{-F}$ to be PSPACE-complete for $k \geq 2$. Previously, $\text{PUSH-}*\text{-F}$ was shown to be PSPACE-complete by Bremner, O’Rourke, and Shermer [BOS94].

On the other hand, Holzer and Schwoon [HS01] proved that the *Atomix* puzzle is PSPACE-complete. Atomix is similar to PUSHPUSH-1 in that only a single block can be pushed at once, and when pushed, it slides its maximal extent. One difference is that the pushing is not made by a robot, but rather by an external agent (the player), so any block can be pushed at any time (unless it is surrounded by other blocks). Another difference is the goal: in Atomix, the blocks are distinguished and must be assembled into a particular pattern.

Our results. This paper resolves the complexity of both $\text{PUSHPUSH-}k$ and $\text{PUSHPUSHPUSH-}k$ by proving them PSPACE-complete for any $k \geq 1$. The reduction reuses certain ingredients from both the construction of Holzer and Schwoon [HS01] and the NP-hardness proof for PUSHPUSH-1 by Demaine et al. [DDO00b].

Organization. The paper is organized as follows. In the next section we outline the reduction from intersection of regular languages given by finite automata. Then in Section 3 we describe the implementation of the gadgets and prove their main properties. Finally, in Section 4 we give our main result and summarize the computational complexity of some pushing-block puzzles.

2 Reduction

Our reduction is from the PSPACE-complete problem of testing for non-empty intersection of regular languages [Gal76,Koz77]. More precisely, we are given a number $n \in \mathbb{N}$, and a sequence A_1, \dots, A_n of finite automata, with $A_i = (Q_i, \Sigma, \delta_i, q_i, F_i)$, for $1 \leq i \leq n$, all using a common alphabet Σ . The goal is to determine whether there is a word in Σ^* that is accepted by all A_i , $1 \leq i \leq n$.

We construct a $\text{PUSHPUSH-}k$ instance that has a solution precisely if there is a word that is accepted by all given automata. The basic idea is to force the player to choose a word that is accepted by all finite automata and then to simulate the transition functions of the automata step-by-step.

The construction is based on the observation that any $(k+1) \times (k+1)$ square group of blocks is immobile by a $\text{PUSHPUSH-}k$ robot, so we can view them as effectively pinned to the plane. In fact, most of the blocks in the constructed

instance are arranged in such immobile groups, only a very few small number of blocks can indeed be pushed by the robot.

An interesting property of $\text{PUSH-PUSH-}k$ – as opposed to, for example, $\text{PUSH-}k$ – is that an obstacle block is sometimes helpful: it can be used to control the movement of other blocks. One can design a key-lock mechanism that requires an obstacle to be pushed to a certain position before the robot can pass. Hence, a truly movable obstacle block is a scarce and important resource in our instances. The robot has to accept some workload to get such a block into the desired position. In particular, this includes the simulation of a finite automaton.

2.1 Gadget Symbols and Properties

We introduce some symbols and basic properties of gadgets; see Fig. 2. Each gadget puts certain restrictions on how a single movable obstacle block can be routed through it. None of the basic gadgets restricts the robot’s movement in any way. We delay describing the actual gadgets until Section 3, when their development is motivated.




Fig. 2. Symbols for main gadgets: (a) one-way gadget, (b) choice gadget, (c) merge gadget, (d) cross gadget, and (e) synchronize gadget.

The main properties of these gadgets are as follows.

One-way gadget: A block can be routed through this gadget in one direction only. One-way gadgets are implicit throughout other gadgets by arrow notation.

Choice gadget: An incoming block can be routed to either outgoing port.

Merge gadget: An incoming block from either port can be routed to the outgoing port.

Cross gadget: Blocks can be routed through horizontally and vertically, but a block cannot take a turn between horizontal and vertical.

Synchronize gadget: Only when two incoming blocks have arrived they can pass through the central channel and be distributed to the two outgoing ports, with exactly one block exiting along each outgoing port.

Recall that the robot is free to pass through the gadgets in any direction. Also, all gadgets can be generalized to an arbitrary number of input and output channels.

2.2 Single Step Transition

Next we describe how to simulate a single step of a finite automaton $M = (Q, \Sigma, \delta, q_0, F)$, where Q is the set of states, Σ is the input alphabet, $\delta : Q \times \Sigma \rightarrow Q$ is the (deterministic) transition function, q_0 is the initial state, and F is the set of final (accepting) states. It is then not much harder to simulate several steps and put together several finite automata. Fig. 3 shows in part (a) the symbolic diagram of the transition simulation of the finite automaton, and in part (b) an example for a particular automaton with two states and two symbols, using the transition function $\delta(q_0, a_0) = q_1$; $\delta(q_0, a_1) = \delta(q_1, a_0) = \delta(q_1, a_1) = q_0$.

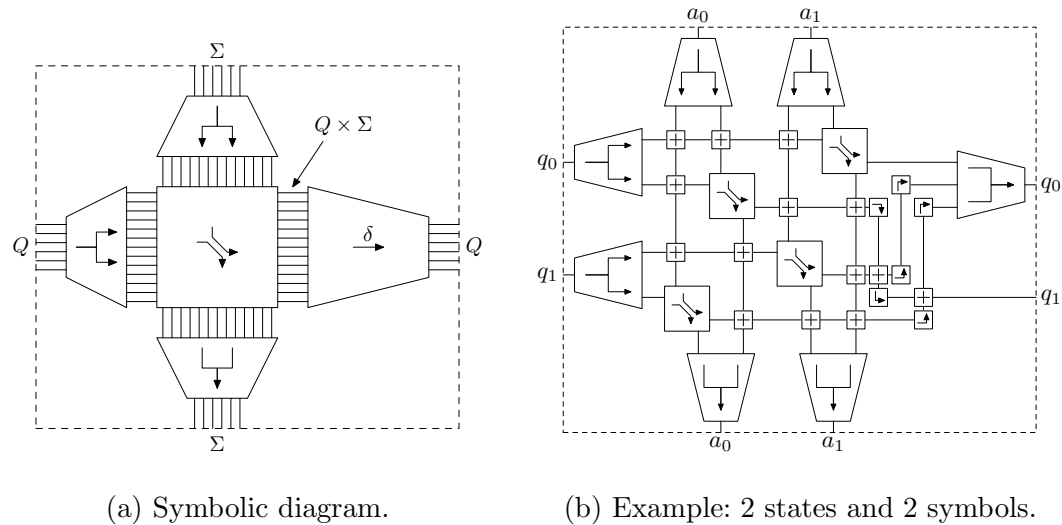


Fig. 3. Transition simulation of an automaton.

On the left of the construction, we suppose that there is a block in exactly one of the state channels $q \in Q$, representing the current state of the finite automaton. On the top of the construction, we suppose that there is a block in exactly one of the symbol channels $a \in \Sigma$, representing the next input symbol.

The idea is that the transition function $\delta : Q \times \Sigma \rightarrow Q$ can be viewed as a matrix, where each row corresponds to a state, and each column corresponds to an input symbol. The construction contains a corresponding $|Q| \times |\Sigma|$ matrix of synchronize gadgets. The current state $q \in Q$ of the automaton determines a particular row in this matrix, and the next input symbol $a \in \Sigma$ determines a particular column, and we want the machine to go to the identified cell (q, a)

of the matrix. Using choice gadgets, we let the state block choose which column it would like to visit, and we let the symbol block choose which row it would like to visit. For the blocks to make any progress beyond the matrix of synchronize gadgets, they must go to a common cell, which is precisely the desired cell (q, a) .

As output, the simulation produces the same symbol a at the bottom of the construction (for later use in another finite automaton) and a new state at the right of the construction, determined by the transition function δ . Specifically, if a block enters the δ gadget on wire $(q, a) \in Q \times \Sigma$, the simulation produces the successor state $\delta(q, a)$. In Figure 3(a), it is hidden that the implementation of the multiwire synchronize gadget uses a number of one-way and cross gadgets, as the concrete example in Figure 3(b) shows.

2.3 Simulation of Several Finite Automata

Now we piece together several instances of Fig. 3 in order to simulate the intersection non-emptiness problem for several finite automata. A schematic view of the construction is given in Fig. 4, where each box labeled “Automaton” corresponds to an instance of Figure 3.

First, the input state channel of each finite automaton contains an initial free block, thus specifying the initial state of the automaton under consideration. Moreover, each input state channel corresponding to a final (accepting) state has an alternate exit route labeled F_i . These exits enable the robot to route a block towards the lock gadget, which is located in the lower-left corner of the figure and is labeled by the word “Key.” This lock corresponds to the automaton termination and accepting the guessed word.

Second, the output state channels loop back to the input state channels of each finite automaton, permeating this information. Third, the symbol channels are strung through all the finite automata, and finally loop back again, with a nondeterministic choice at the very top; this is where the input word is generated, symbol-by-symbol.

Finally, there are gadgets responsible for the robot’s start and finish. Initially, the robot is located at the place marked with the large S. The sequence of locks prevent the robot from immediately accessing its goal position that is labeled with a large X. Because the symbol channels are all strung together, in fact every automaton must terminate and accept at the same time. Thus, one block per automaton can reach the “finish,” unlocking the corresponding lock, precisely if there is a common word accepted by all of the finite automata.




Fig. 4. Non-empty intersection simulation for several finite automata.

3 Gadgets

This section details the **PUSH** gadgets used in the reduction of the previous section. We begin with some basic observations and conventions, following [DDO00b].

First, observe that any $(k + 1) \times (k + 1)$ group of blocks is immobile by a $\text{PUSHPUSH-}k$ robot, so we can view them as effectively pinned to the plane. Such groups of blocks are shaded dark in our figures. Note that this shading is merely an aid to structure the drawings visually; there is no conceptual difference between light and dark blocks. We use such immobile groups to build *channels*: a channel is a $1 \times x$ or $x \times 1$ group of empty cells surrounded on both sides by two $(k + 1) \times x$ or $x \times (k + 1)$ groups of pinned blocks.

In all but the synchronize gadget, our constructions will use these width-1 channels connected at orthogonal junctions of degree between two and four. We can view such a construction as an orthogonal graph, where the edges represent channels and the vertices represent orthogonal junctions. We use circles to denote movable blocks in the channels or at the channel junctions.

For the remainder of this section, we will restrict our attention to PUSHPUSH-1 . It is straightforward to extend the gadgets to the general $\text{PUSHPUSH-}k$ constraint, in most cases – to be precise, except for the lock gadget – simply by thickening the channel insulation. Note that even for $k > 1$ there is only one symbol block plus one state block for each automaton to be pushed around. In fact, the construction is done in such a way that the robot will never be able to make use of its ability to push several blocks simultaneously.

3.1 One-way, Choice, Merge, and Cross Gadgets

We start with the gadgets depicted in Fig. 2 save the synchronize gadget. For each gadget, Fig. 5 shows the orthogonal-graph view on the top and an actual PUSHPUSH-1 puzzle on the bottom.

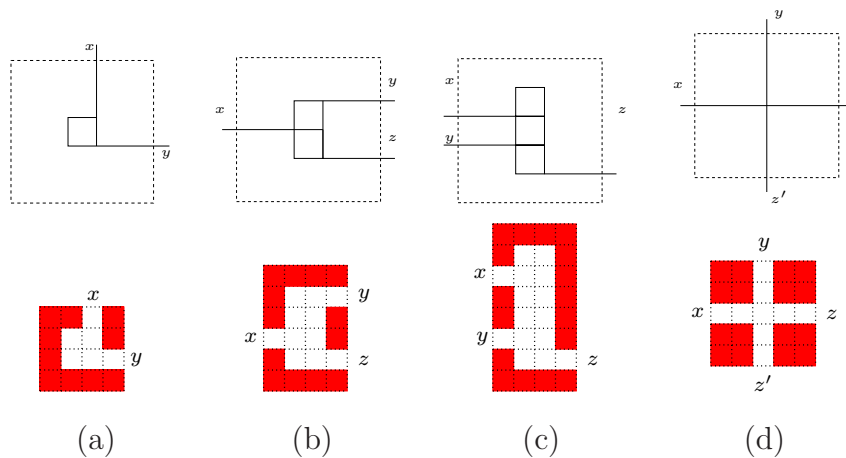


Fig. 5. Gadget implementations: (a) one-way, (b) choice, (c) merge, and (d) cross gadget.

These gadgets have the following obvious properties, which we state without

proof. Except for the catalyst/synchronize gadget and the lock gadget, we assume that only one movable block resides in a particular gadget at any time. This property is fulfilled in our construction.

Lemma 1 *In the one-way, choice, merge, corner, and cross gadget, the robot is free to travel in any direction. For blocks, the following properties hold, assuming that at most one block occupies the gadget at once:*

- (1) *In the one-way gadget, the robot may move a block from x to y , but not from y to x .*
- (2) *In the choice gadget, the robot may move a block only from x to y , or from x to z .*
- (3) *In the merge gadget, the robot may move a block only from x to z , or from y to z .*
- (4) *In the cross gadget, the robot may move a block only from x to z or y to z' , or vice versa.*

3.2 Synchronize Gadget

We next detail the synchronize gadget, which is built out of two *catalyst chambers* wired in sequence as illustrated in Fig. 6(a). The difference between catalyst chamber and synchronize gadget is that a catalyst chamber has no “central channel”, so the two blocks are not required to leave the gadget simultaneously. The second catalyst gadget acts as an “after-burner”, ensuring that both blocks leave the first gadget. The catalyst gadget is illustrated in Fig. 6(b–f).

Lemma 2 *In the catalyst gadget,*

- (1) *if the robot moves one block into the gadget alone from either x or y , it cannot move the block to either z or w ;*
- (2) *if the robot moves two blocks into the gadget, there is a sequence of moves such that both blocks can leave the gadget through different output channels; and*
- (3) *if the robot moves two blocks into the gadget, the blocks cannot both be pushed through the same output channel z or w .*

The robot is free to travel in any direction.

PROOF. Property 1 is easy to verify; once a block is moved into the gadget alone, it can only be pushed around cyclically inside the gadget, but never enter one of the output channels. Property 2 is shown by the example sequence shown in Fig. 6, parts (b) to (d). Finally, for Property 3, note that the blocks can basically only be pushed around in the cycle C that is shaded grey in

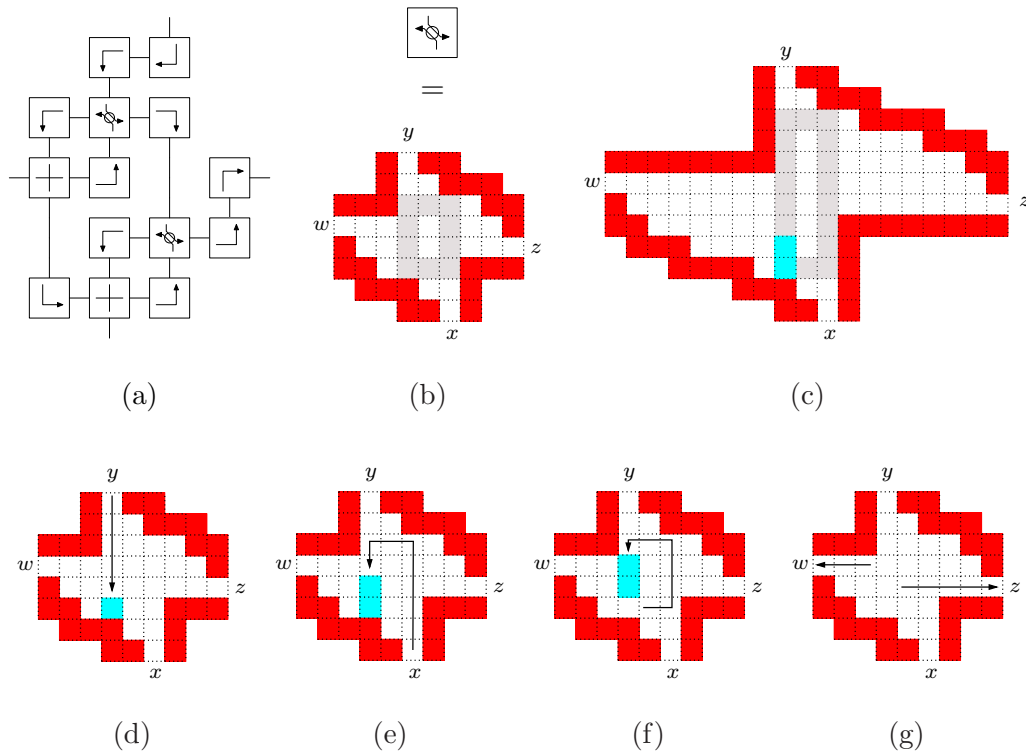


Fig. 6. (a) Synchronize gadget, constructed in terms of (b) catalyst gadget, (c) a catalyst gadget for PUSH-PUSH-PUSH-3, and (d)–(g) the traversal sequence of (b).

Fig. 6(b). As soon as the first block is pushed out from C , it either leaves the gadget via an output channel, leaves the gadget via an input channel, or it gets stuck in a corner forever. In the latter two cases, we are done. In the former case, which can only occur if both block are placed immediately adjacent on C , we are left with one block on C . Either this block is immediately be pushed out, or it will be stuck in C as in Property 1. This concludes the proof, because there are no two neighboring positions on the cycle from which blocks can be pushed out to the same channels. \square

No changes are required for the catalyst gadget to work with the PUSH-PUSH- k constraint for $k > 1$. However, for PUSH-PUSH-PUSH- k with $k > 1$, the operations shown in Fig. 6 (e)–(f) do not lead to (f) but back to (e) instead, since the $2 \leq k$ blocks continue to move down. Therefore the gadget has to be “blown up” such that an sequence analogous to the one in Fig. 6 ends up with a sequence of $k + 1$ blocks stacked up in the middle. For an arbitrary $k \in \mathbb{N}$ the catalyst gadget for PUSH-PUSH-PUSH- k has an interior cycle C of height $2(k + 1)$ and width three, and the sizes of the staircase steps on its boundary decrease gradually from k to one when going from x towards w or z towards y , respectively. Moreover, the gadget initially contains $k - 1$ movable blocks, stacked up in a corner of C . The instance for PUSH-PUSH-PUSH-3 is depicted

in Fig. 6(c).

3.3 Lock Gadget

Finally, we describe the lock gadget. It prevents the robot from reaching its final destination unless it has enough blocks to unlock the gadget. A single lock gadget is shown in Fig. 7(a). By cascading several of these gadgets in sequence, we obtain a generalized lock gadget with several key input channels, each requiring a key.




Fig. 7. Lock gadget.

The lock is quite different compared to the gadgets we have seen so far in that it restricts the robot's movement. There are a number of movable blocks which may change position whenever the robot traverses a lock. That is, a lock has an associated state. We distinguish between three different states: a lock as shown in Fig. 7(a) is *closed*, if a key block has been pushed in from y (see Fig. 7(b)), the lock is *open*, and all other possible states are referred to as *broken*. The possible traversals of a lock gadget and the resulting changes in state are summarized by the following lemma.

Lemma 3 *The robot may pass through a lock gadget from x to z if and only if the lock is open. The robot may pass through a lock gadget from x to y if and only if the lock is closed. After any of these traversals the lock is broken. The robot may change the state of a lock from locked to open by pushing in a block from y . Except for the traversals described above, there is no way for the robot to pass through the gadget between any two distinct ports $(\{x, y, z\})$ in any state.*

PROOF. First consider a locked gadget: obviously, it cannot be traversed from either y or z . In fact, Block 0 is the only block that can be pushed at all. If it is pushed down, no further block movement is possible in the gadget. If Block 0 is pushed right, Block 1 can only be pushed up afterwards such that it gets stuck in the corner and blocks the connection to y or z forever. Hence, there is only one way to proceed: pushing Block 0 left and thereby blocking port x forever (including the way back). Moreover, the only way to get out of the gadget is to proceed as in Figure 7(c), except that there is no Block 8 such that Block 4 moves all the way until it hits Block 5, thereby forever blocking port z as well.

Now consider an open lock: again, the gadget cannot be traversed from either y or z . As above, the only way to traverse the lock from x is shown in Figure 7(c). Afterwards, both x and y are blocked forever. \square

For $\text{PUSHPUSH-}k$ with $k > 2$ a slight modification of the lock gadget is required on top of the usual thickening: Block 4 has to be replaced by a sequence of k blocks. Basically, we copy the column to the right of Block 4 and insert another block to the right of Block 4 (k times). In the proof of Lemma 3 there is a minor difference as well, since for $k > 2$ the robot might first push Block 0 to the right without immediately blocking the whole gadget. To avoid introducing additional states, we just place Block 0 and Block 1 next to each other from the very beginning. Fig. 7(d) shows the resulting gadget for $\text{PUSHPUSH-}2$. Finally observe that the lock gadget immediately extends to $\text{PUSHPUSH-}k$.

This completes the description of the gadgets.

4 Conclusions

By the construction presented in Section 2 and by the obvious fact that $\text{PUSH-}k$ belongs to PSPACE , we conclude the following:

Theorem 4 *$\text{PUSHPUSH-}k$ and $\text{PUSHPUSH-}k$ are PSPACE-complete for any $k \in \mathbb{N}$.*

Finally, we summarize some of the computational complexity results on pushing-block puzzles in Tab. 8. The main open problem that remains is to resolve the complexity of PUSH- k and PUSH- $*$ between NP and PSPACE.

Problem	Parameter		
	-1	- k	-*
PUSH	NP-hard	NP-hard	NP-hard
PUSHPUSH	PSPACE-compl.	PSPACE-compl.	NP-hard
PUSH-X	NP-compl.	NP-compl.	NP-compl.
PUSH-F	NP-hard	PSPACE-compl.	PSPACE-compl.

Tab. 8. Computational complexity of some block sliding puzzles.

References

- [BOS94] David Bremner, Joseph O’Rourke, and Thomas C. Shermer. Motion planning amidst movable square blocks is PSPACE complete. Unpublished Draft, 1994.
- [DDO00a] Erik D. Demaine, Martin L. Demaine, and Joseph O’Rourke. PushPush and Push-1 are NP-hard in 2D. In *Proceedings of the 12th Annual Canadian Conference on Computational Geometry*, pages 211–219, Fredericton, Canada, August 2000. <http://www.cs.unb.ca/conf/cccg/eProceedings/26.ps.gz>.
- [DDO00b] Erik D. Demaine, Martin L. Demaine, and Joseph O’Rourke. PushPush is NP-hard in 2D. Technical Report 065, Department of Computer Science, Smith College, Northampton, MA, January 2000. <http://arXiv.org/abs/cs.CG/0001019>.
- [Dem01] Erik D. Demaine. Playing games with algorithms: Algorithmic combinatorial game theory. In *Proceedings of the 26th Symposium on Mathematical Foundations in Computer Science*, volume 2136 of *Lecture Notes in Computer Science*, pages 18–32, Mariánské Lázně, Czech Republic, August 2001. Full paper available at <http://www.arXiv.org/abs/cs.CC/0106019>.
- [DH01] Erik D. Demaine and Michael Hoffmann. Pushing blocks is NP-complete for noncrossing solution paths. In *Proceedings of the 13th Canadian Conference on Computational Geometry*, pages 65–68, Waterloo, Canada, August 2001. <http://compgeo.math.uwaterloo.ca/~cccg01/proceedings/long/eddemaine-24711.ps>.

- [DHH02] Erik D. Demaine and Michael Hoffmann. Push-2F is PSPACE-Complete. In *Proceedings of the 14th Canadian Conference on Computational Geometry*, pages 31–35, Lethbridge, Canada, August 2002. <http://www.inf.ethz.ch/~hoffmann/pub/dhh-ppc-02.ps>.
- [Gal76] Zvi Galil. Hierachies of complete problems. *Acta Informatica*, 6(1):77–88, 1976.
- [H00] Michael Hoffmann. Push-* is NP-hard. In *Proceedings of the 12th Annual Canadian Conference on Computational Geometry*, pages 205–210, Fredericton, Canada, August 2000. <http://www.cs.unb.ca/conf/cccg/eProceedings/13.ps.gz>.
- [HS01] Markus Holzer and Stefan Schwoon. Assembling molecules in Atomix is hard. Technical Report TUM-I0101, Technische Universität München, May 2001.
- [Koz77] Dexter Kozen. Lower bounds for natural proof systems. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, pages 254–266, Providence, RI, 1977.
- [OS99] Joseph O’Rourke and the Smith Problem Solving Group. PushPush is NP-hard in 3D. Technical Report 064, Department of Computer Science, Smith College, Northampton, MA, November 1999. <http://arXiv.org/abs/cs/9911013>.