

1 Computational Complexity of Motion Planning of 2 a Robot through Simple Gadgets

3 **Erik D. Demaine**

4 MIT CSAIL, 32 Vassar Street, Cambridge, MA 02139, USA
5 edemaine@mit.edu

6 **Isaac Grosz**¹

7 MIT CSAIL, 32 Vassar Street, Cambridge, MA 02139, USA
8 isaacg@alum.mit.edu

9 **Jayson Lynch**

10 MIT CSAIL, 32 Vassar Street, Cambridge, MA 02139, USA
11 jaysonl@mit.edu

12 **Mikhail Rudoy**²

13 MIT CSAIL, 32 Vassar Street, Cambridge, MA 02139, USA
14 mrudoy@gmail.com

15 — Abstract —

16 We initiate a general theory for analyzing the complexity of motion planning of a single robot
17 through a graph of “gadgets”, each with their own state, set of locations, and allowed traversals
18 between locations that can depend on and change the state. This type of setup is common
19 to many robot motion planning hardness proofs. We characterize the complexity for a natural
20 simple case: each gadget connects up to four locations in a perfect matching (but each direction
21 can be traversable or not in the current state), has one or two states, every gadget traversal is
22 immediately undoable, and that gadget locations are connected by an always-traversable forest,
23 possibly restricted to avoid crossings in the plane. Specifically, we show that any single nontrivial
24 four-location two-state gadget type is enough for motion planning to become PSPACE-complete,
25 while any set of simpler gadgets (effectively two-location or one-state) has a polynomial-time
26 motion planning algorithm. As a sample application, our results show that motion planning
27 games with “spinners” are PSPACE-complete, establishing a new hard aspect of *Zelda: Oracle*
28 *of Seasons*.

29 **2012 ACM Subject Classification** Theory of computation → Problems, reductions and com-
30 pleteness

31 **Keywords and phrases** PSPACE, hardness, motion planning, puzzles

32 **Digital Object Identifier** 10.4230/LIPIcs.FUN.2018.18

33 **1 Introduction**

34 Many hardness proofs are based on *gadgets* — local pieces, each often representing corre-
35 sponding pieces of the input instance, that combine to form the overall reduction. Garey and
36 Johnson [7] called gadgets “basic units” and the overall technique “local replacement proofs”.
37 The search for a hardness reduction usually starts by experimenting with small candidate

¹ Now at Carnegie Mellon University.

² Now at Google Inc.



38 gadgets, seeing how they behave, and repeating until amassing a sufficient collection of
 39 gadgets to prove hardness.

40 This approach leads to a natural question: what gadget sets suffice to prove hardness?
 41 There are many possible answers to this question, depending on the precise meaning of
 42 “gadget” and the style of problem considered. Schaefer [11] characterized the complexity of
 43 all “Boolean constraint satisfiability” gadgets, including easy problems (2SAT, Horn SAT,
 44 dual-Horn SAT, XOR SAT) and hard problems (3SAT, 1-in-3SAT, NAE 3SAT). Constraint
 45 Logic [8] proves sufficiency of small sets of gadgets on directed graphs that always satisfy one
 46 local rule (weighted in-degree at least 2), in many game types (1-player, 2-player, 2-team,
 47 polynomially bounded, unbounded), although the exact minimal sets of required gadgets
 48 remain unknown. Both of these general techniques naturally model “global” moves that can
 49 be made anywhere at any time (while satisfying the constraints). Nonetheless, the techniques
 50 have been successful at proving hardness for problems where moves must be made local to
 51 an agent/robot that traverses the instance.


52 In this paper, we introduce a general model of gadgets that naturally arises from *single-*
 53 *agent motion planning problems*, where a single agent/robot traverses a given environment
 54 from a given start location to a given goal location. Our model is motivated by the plethora
 55 of existing hardness proofs for such problems, such as Push-1, Push-*, PushPush, and
 56 Push-X [3]; Push-2-F [5]; Push-1 Pull-1 [4,9]; as well as several Nintendo video games studied
 57 at recent FUN conferences [1,6].

58 1.1 Gadget model

59 In general, we model a *gadget* as consisting of one or more *locations* (entrances/exits) and one
 60 or more *states*. (In this paper, we will focus on gadgets with at most two states.) Each state
 61 s of the gadget defines a labeled directed graph on the locations, where a directed edge (a, b)
 62 with label s' means that the robot can enter the gadget at location a and exit at location b ,
 63 and that such a traversal forcibly changes the state of the gadget to s' . Equivalently, a gadget
 64 is specified by its *state space*, a directed graph whose vertices are state/location pairs, where
 65 a directed edge from (s, a) to (s', b) represents that the robot can traverse the gadget from a
 66 to b if it is in state s , and that such traversal will change the gadget’s state to s' . Gadgets
 67 are *local* in the sense that traversing a gadget does not change the state of any other gadgets.

68 A *system of gadgets* consists of gadgets, their initial states, and *connections* between
 69 disjoint pairs of locations (forming a matching). If two locations a, b of two gadgets (or the
 70 same gadget) are connected, then the robot can traverse freely between a and b (outside the
 71 gadgets). (Equivalently, we can think of locations a and b as being identified.) These are
 72 all the ways that the robot can move: exterior to gadgets using connections, and traversing
 73 gadgets according to their current states. In a *puzzle*, we are given a system of gadgets, the
 74 robot starts at a specified start location, and we want to find a sequence of moves that brings
 75 the robot to a specified goal location. The main problem we consider here is the obvious
 76 decision problem: is the given puzzle solvable?

77 One type of gadget we always allow in this paper is the **branch-**
 78 **ing hallway** gadget, which has one state and three locations, and
 79 always allows traversal between all pairs of locations; see Figure 1.
 80 In other words, upon reaching such a gadget, the robot is free to
 81 choose and move to any of the three locations. Connecting together
 82 multiple branching hallways allows us to effectively connect the other
 83 gadgets’ locations according to an arbitrary forest (as described in
 84 the abstract).



■ **Figure 1** Branching hallway gadget

85 All other gadgets we consider in this paper are “deterministic” and “reversible”. A gadget
86 is *deterministic* if its state space has maximum out-degree ≤ 1 , i.e., a robot entering the
87 gadget at some location a in some state s (if possible) can exit at only one location b and
88 one new state s' . A gadget is *reversible* if its state space has the reverse of every edge, i.e., it
89 is the bidirectional version of an undirected graph. Thus a robot can immediately undo any
90 gadget traversal.³ Together, determinism and reversibility are equivalent to requiring that
91 the state space is the bidirectional version of a matching.

92 Other than the (one-state) branching hallway, we further require that the states of a
93 gadget differ only in their orientations of the possible traversals. More precisely, a k -tunnel
94 gadget has $2k$ locations, paired in a perfect matching whose pairs are called *tunnels*, such
95 that each state defines which direction or directions each tunnel can be traversed.

96 We also consider *planar* systems of gadgets, where the gadgets and connections are drawn
97 in the plane without crossings. Planar gadgets are drawn as small regions (say, disks) with
98 their locations as points in a fixed clockwise order along their boundary. A single gadget type
99 thus corresponds to multiple planar gadget types, depending on the choice of the clockwise
100 order of locations. Connections are drawn as paths connecting the points corresponding to
101 the endpoint locations, without crossing gadget interiors or other connections.

102 1.2 Our results

103 We characterize the computational complexity of deciding puzzle solvability when the allowed
104 gadgets consist of the branching hallway and any number of deterministic reversible ≤ 2 -state
105 k -tunnel gadgets, for any k . Specifically, if there is at least one gadget type that is not
106 equivalent to a 1-state or 1-tunnel gadget, then the problem is PSPACE-complete; and
107 otherwise, the problem is in P. The same characterization holds for planar systems of gadgets;
108 thus, in applications, we do not have to worry about building a crossover gadget (which is
109 often the most difficult).

110 In Section 3, we sketch our proof from [4] that motion planning with two-toggle-locks
111 and crossovers is PSPACE-complete. In Section 4, we prove that one particular gadget,
112 the antiparallel two-toggle, can simulate a variety of other gadgets, eventually including a
113 two-toggle-lock and a crossover. As a consequence, motion planning with the antiparallel
114 two-toggle is PSPACE-complete. In Section 5, we show that all nontrivial deterministic
115 reversible 2-state, 2-tunnel gadgets can simulate the antiparallel two-toggle. As a consequence,
116 each corresponding motion planning problem is PSPACE-complete. In Section 7, we extend
117 these results to give a precise hardness characterization for the motion planning problem
118 with each deterministic reversible 2-state k -tunnel gadget.

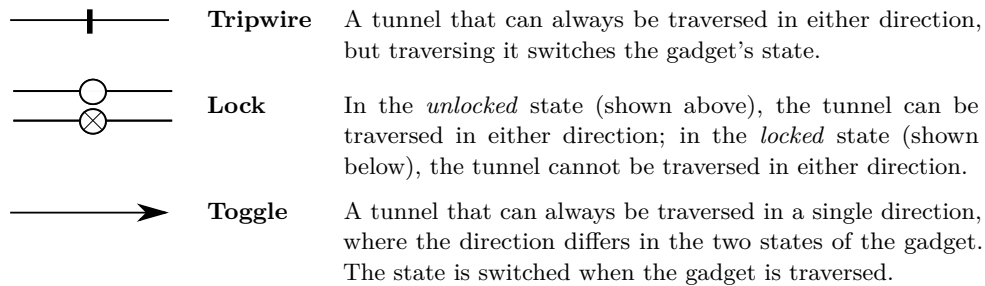
119 We also partially characterize the computational complexity of deterministic reversible
120 ≤ 2 -state gadgets with three locations. In particular, we study spinners and deterministic
121 forks, as described in Section 6.

122 We hope that our approach will be useful for establishing hardness of many real-world
123 motion planning problems and puzzles. As a sample application, our results allow us to
124 establish a new PSPACE-hard aspect of the Nintendo video game *Zelda: Oracle of Seasons*
125 (which features spinners) Section 6.

³ This notion is different than the sense of “reversible” in reversible computing, which would mean that we could derive which move to undo from the current state.

126 **2 Gadget Basics**


127 To categorize the possible deterministic reversible 2-state 2-tunnel gadget types, we first
 128 categorize the possible tunnel types in such a gadget. A tunnel is *trivial* if it is either never
 129 traversable or always traversable. A trivial tunnel can always be split into a separate 1-state
 130 1-tunnel gadget, so we can ignore them. What remain are three possible *nontrivial* tunnel
 131 types:



133 There are six ways to combine these tunnel types into pairs. Two combinations, Lock–Lock
 134 and Tripwire–Tripwire, are trivial combinations equivalent to one-state gadgets in which
 135 each tunnel is either always traversable in both directions or never traversable. Thus we
 136 restrict our attention to the four other combinations, listed below. Because we are interested
 137 in planar systems, we consider the multiple planar gadgets for each nontrivial combination.
 138 (We do, however, treat a gadget and its reflection as equivalent.) As a result, there are nine
 139 different nontrivial two-tunnel two-state gadgets, abbreviated and listed below. The bulk of
 140 our paper focuses on the six gadgets shown in Figure 2, which omits most crossing variants.

- 141 **1. Tripwire–Lock:** Traversing the tripwire makes the other tunnel flip between being
 142 passable and impassable, causing it to ‘lock’ or ‘unlock’. There are crossing and non-
 143 crossing varieties, abbreviated **CWL** (crossing wire lock) and **NWL** (non-crossing wire
 144 lock).
- 145 **2. Toggle–Lock:** Traversing the toggle flips the lock tunnel between being passable and
 146 impassable. Crossing the lock tunnel, by definition, does not change the state of the
 147 gadget. Notice that one direction of the toggle corresponds to an open lock and the other
 148 direction to the closed lock. There are crossing and non-crossing varieties, abbreviated
 149 **CTL** (crossing toggle lock) and **NTL** (non-crossing toggle lock).
- 150 **3. Tripwire–Toggle:** Here traversing either the tripwire or the toggle flips the direction of
 151 the toggle. There are crossing and non-crossing varieties, abbreviated **CWT** (crossing
 152 wire toggle) and **NWT** (non-crossing wire toggle).
- 153 **4. Toggle–Toggle:** Also known as a **2-toggle** [4]. Traversing either toggle flips the direction
 154 of both of them. This is the only case where there are two directed tunnels, leading
 155 to three possibilities: crossing, parallel, and anti-parallel. They are abbreviated **C2T**
 156 (crossing 2-toggle), **P2T** (parallel 2-toggle), and **AP2T** (anti-parallel 2-toggle).

157 In this paper we will often need to discuss putting gadgets together to create new behavior.
 158 We will do so by creating a system of gadgets that is “equivalent” to some target gadget,
 159 thereby “simulating” that gadget. Two systems of gadgets are *equivalent* if there is a bijective
 160 correspondence between their locations and a correspondence between their states such
 161 that the allowed transitions for all (locations, state) pairs are the same under these two
 162 correspondences. We will say that a gadget or set of gadgets *simulates* a target gadget if it is



■ **Figure 2** Six of the nine deterministic reversible 2-state gadgets on two tunnels. We leave out the CWL, CTL, and CWT gadgets as they are not heavily used in the paper.

163 possible to combine gadgets from the set (possibly using duplicates) such that the resulting
 164 system is equivalent to the target gadget. We will always implicitly allow the use of the
 165 branching hallway gadget in these constructions. In all cases, these constructions will be
 166 planar.

167 2.1 Closure Properties

168 ► **Lemma 2.1.** *Any system of gadgets composed of two reversible gadgets is reversible.*

169 **Proof.** Consider any transition through the system formed by composing two reversible
 170 gadgets. This transition is a walk through the gadgets and connections that form a system.
 171 Since both gadgets are reversible, it is possible for the robot to enact the exact reverse of
 172 this walk after the walk is done. This will exactly reverse the effect of the walk within each
 173 gadget. Thus, it is possible to reverse the entire transition.

174 Since every transition of the system can be reversed, the system is reversible. ◀

175 Since all of the gadgets we consider in this paper are reversible, Lemma 2.1 means our
 176 systems will all be reversible as well.

177 ► **Lemma 2.2.** *Any system of gadgets composed of two deterministic reversible gadgets is
 178 deterministic and reversible.*

179 **Proof.** The state space of a reversible, deterministic gadget is an undirected matching of
 180 some (state, location) pairs to each other. This is a necessary and sufficient characterization of
 181 reversible, deterministic gadgets.

182 When we compose two such gadgets, we create paths through the pair of gadgets. However,
 183 no (state, location) pair has more than two edges: One connection to the other gadget, and
 184 one edge through its original gadget. Moreover, any (state, location) pair that forms an
 185 external location has at most one edge, as it does not connect to the other gadget. As a
 186 consequence, the path from any external location through the gadget is either a deterministic
 187 path to another external location, or a dead end. There is no branching, as branching would
 188 require a location with three edges.

189 Thus, the resultant object is deterministic. By Lemma 2.1 it is reversible as well. ◀

190 2.2 PSPACE Membership

191 ► **Lemma 2.3.** *Deciding puzzle solvability is in PSPACE.*

192 **Proof.** The entire state of the system can be described by the current state of the gadgets
 193 and the location of the agent. The gadgets have a polynomial number of states and there
 194 can only be a polynomial number of gadgets. Since the entire state of the board fits in a

195 polynomial amount of space, we can non-deterministically search for a solution, showing
 196 containment in NPSPACE. Savich's Theorem [10] gives $PSPACE = NPSPACE$. ◀

197 **3 2-toggle-lock and crossover motion planning is PSPACE-complete**

198 In [4] we showed that motion planning with 4-toggles and crossovers is PSPACE-complete.
 199 In that construction, the crucial gadget turned out to be a 2-toggle-lock, which is a 3-tunnel,
 200 2-state gadget with two locks and a tunnel. The 4-toggle was not used in any way after the
 201 construction of the 2-toggle-lock, showing that 2-toggle-locks and crossovers are PSPACE-
 202 hard. For convenience we sketch the proof, with some refinement. One should refer to the
 203 prior paper for a more detailed and rigorous proof.

204 ▶ **Definition 3.1.** 3QSAT is the following decision problem. Given a fully quantified boolean
 205 formula in prenex normal form and in conjunctive normal form with no more than three
 206 variables per clause, decide whether the formula is true.

207 ▶ **Theorem 3.2.** *Motion planning with 2-toggle-locks and crossovers is PSPACE-hard.*

208 We reduce from 3QSAT to motion-planning with 2-toggle-locks and crossovers. To do so we
 209 need to construct clauses, universal variables, and existential variables. Literals will consist
 210 of a 2-toggle-lock which will be set from the 2-toggle side and checked by passing through the
 211 lock. Clauses are composed of a branching hallway that leads through each of its associated
 212 literals.

213 Existential variables will be a branching hall with a group of toggle-locks in series. Passing
 214 through in one direction opens the locks of the gadgets representing true literals of that
 215 variable while closing the locks of the false ones. Going through the other way allows this to
 216 be undone, as the system is reversible.

217 To construct universal quantifiers we connect up the 2-toggle sections as in Figure 3,
 218 where each universal gadget consists of several antiparallel 2-toggles with locks. Each of these
 219 gadgets sends the robot forward in one state or back to the beginning in the other state, and
 220 flips the state. Repeatedly entering from the left iterates through all configurations of the
 221 states, so the robot must check all of the possible values for the universal variables. The goal
 222 state lies at the far end of the eries of universal gadgets.

223 For both the existentials and the universals, the variables are actually a long series of
 224 2-toggle-locks with one lock for each literal of the variable in the formula.

225 When putting this all together, as in Figure 3, we need to ensure that the robot cannot
 226 sneak back into the variable gadget and change existential settings it shouldn't be allowed
 227 to access, namely those existentials beyond the universal it just emerged from. To do this
 228 we construct a simple system that puts a lock on the return pathway at the end of each
 229 universal variable which only allows passage if the prior variable is set to false. Since the
 230 robot will have just exited from a variable which was set to true, this prevents the robot
 231 from moving forward in the variable chain. In addition, all earlier variables are false allowing
 232 the robot to travel back to the formula, since the universal gadgets take on incrementing
 233 binary values with each loop through the gadget. Since those existential variables are ones
 234 the robot was allowed to set to any value on the prior passage, going back and changing
 235 them now gives no advantage over having set them to that value earlier.

236 This safeguard is the one difference from the prior construction, which checked the values
 237 of all prior universal variables, requiring a quadratic blow-up in number of gadgets. The need
 238 for crossovers and a 2D layout will still create a quadratic blowup in problem size overall,

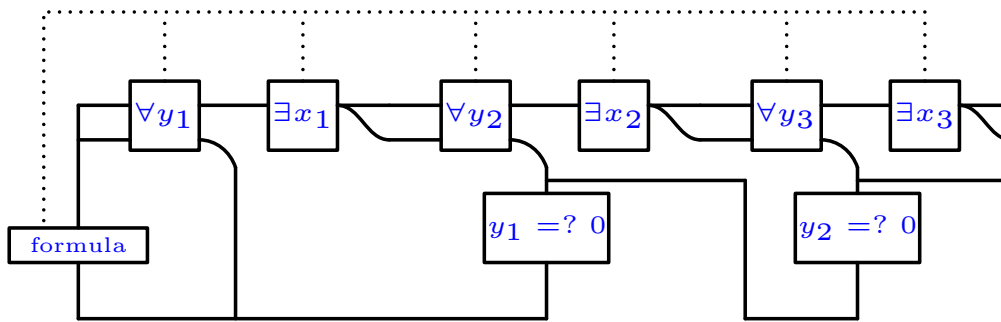


Figure 3 Structure of the QSAT reduction.

but this simplification seemed worth noting and should allow for the 3D result to cause only a linear blowup in problem size.

With this guard in place, the robot can only reach the goal state by demonstrating a solution to the 3QSAT instance, after iterating through all settings of the universal gadget. ◀

4 Antiparallel 2-toggle motion planning is PSPACE-complete

We will show that the question of whether a robot in a system of antiparallel 2-toggle gadgets can reach a specified goal location is PSPACE-complete. To do so, we will simulate various other gadgets using AP2T gadgets, eventually simulating 2-toggle-locks and crossovers. Since motion planning with 2-toggle-locks and crossovers is PSPACE-complete, this implies that AP2T motion planning is PSPACE-complete.

► **Theorem 4.1.** *Motion Planning with AP2T gadgets is PSPACE-complete.*


We will simulate the gadgets needed for the PSPACE-completeness proof, and a wide variety of other intermediate gadgets to help us get there. The steps are as follows:

1. Simulate a C2T, using AP2Ts. Lemma 4.2.
2. Simulate a P2T, using C2Ts. Lemma 4.3.
3. Simulate a NTL, using AP2Ts, C2Ts and P2Ts. Lemma 4.4.
4. Simulate various types of 2-toggle locks, with “round” and “stacked” internal connections. The types of internal connections are described in Section 4.1, and the constructions are given in Lemmas 4.6 and 4.7.
5. Simulate a NWL, using the stacked antiparallel 2-toggle lock. Lemma 4.8.
6. Simulate a stacked tripwire-lock-tripwire, using NWLs. Lemma 4.9
7. Simulate a crossover, using stacked tripwire-lock-tripwires. Lemma 4.10


With a 2-toggle lock and a crossover constructed, we can apply Theorem 3.2 to show that motion planning with AP2Ts is PSPACE-hard. Adding in Lemma 2.3, we find that it is PSPACE-complete. ◀

► **Lemma 4.2.** *Antiparallel 2-toggles (AP2Ts) simulate a crossing 2-toggle (C2T).*


Proof. The construction is given in Figure 4. In the state of the construction shown in the figure, there are two possible transitions: the robot can move from the upper left to the bottom right of the construction, or from the upper right to the bottom left. Either of those transitions toggles both AP2Ts, leaving the construction mirrored top to bottom. Thus, the construction has two states. The possible traversals in one state (as shown above) are from



■ **Figure 4** Anti-parallel 2-toggles simulate a crossing 2-toggle



■ **Figure 5** Crossing 2-toggles simulate a parallel 2-toggle



■ **Figure 6** 2-toggles simulate 1-toggle-lock.

270 the top left to the bottom right and from the top right to the bottom left, while the possible
 271 traversals in the other state are (by symmetry) from the bottom left to the top right and
 272 from the bottom right to the top left. Following any of these traversals swaps the state of
 273 the construction. Notice that this is exactly the behavior of a C2T.

274 If the robot enters the construction shown from the upper left, upon reaching the center
 275 the robot can only proceed to the bottom right, or come back the way it came. Therefore, the
 276 upper left to bottom right transition is the only possible transition from that location. By
 277 symmetry, the same is true from top left to bottom right. Thus, the one traversal described
 278 for each location in each state is the only one possible. ◀

279 ▶ **Lemma 4.3.** *Crossing 2-toggles (C2Ts) simulate a parallel 2-toggle (P2T).*

280 **Proof.** The construction is given in Figure 5. In the state of the construction shown in
 281 the figure, there are two possible transitions: the robot can move from the top left to the
 282 top right of the construction, or from the bottom left to the bottom right. Either of these
 283 transitions toggles both C2Ts, leaving the construction mirrored left to right. The allowed
 284 traversals in one state (as shown above) are from the top left to the top right and from
 285 the bottom left to the bottom right, while the allowed traversals in the other state (by
 286 symmetry) from the top right to the top left and from the bottom right to the bottom left.
 287 Following any of these traversals swaps the state of the construction. Notice that this is
 288 exactly the behavior of a P2T.

289 Since the system is composed entirely of C2Ts (without even branching hallways), which
 290 are both reversible and deterministic, the result is also both reversible and deterministic, by
 291 Lemma 2.2. Thus, the one transition described for each location in each state is the only
 292 transition possible. ◀

293 ▶ **Lemma 4.4.** *2-toggles (AP2Ts, P2Ts and C2Ts) simulate a noncrossing toggle lock (NTL).*

294 **Proof.** The construction is shown in Figure 6.

295 In this lemma, we will refer to toggles 1 and 2 in the figure as the “outer toggles”, toggles
 296 3 and 4 as the “middle toggles”, and toggles 5 and 6 as the “bottom toggles”. We will call
 297 the pathway through the lower tunnels of the bottom toggles the “bottom tunnel” of the
 298 overall gadget, and the rest of the gadget the “middle tunnel” of the overall gadget.

299 An NTL has two externally observable states: locked, and unlocked. The locked state
 300 corresponds to the upper tunnels of the bottom toggles oriented out, and the unlocked state
 301 corresponds to the bottom toggles oriented in. The unlocked state is shown in Figure 6.

302 In this gadget, there are two internal states corresponding to each external state: with
 303 the horizontal tunnels of the middle toggles both oriented left, and with both oriented right.
 304 The only accessible states of this gadget are the states with the outer toggles oriented in, the
 305 middle toggles oriented both left or both right, and upper pathways of the bottom toggles
 306 oriented both in or both out. We will show that the gadget allows exactly the traversals of
 307 the NTL from these configurations, and cannot be left in any other configuration.

308 The bottom tunnel traversals are straightforward — the bottom tunnel acts as a toggle,
 309 and a traversal flips both bottom toggles, and hence the externally observable state.

310 Also clearly, the robot cannot move between the bottom tunnel and the middle tunnel.

311 Now, we wish to establish that in the unlocked state, the robot can always traverse the
 312 middle tunnel in either direction. In the state shown, the middle tunnel may be traversed
 313 from external location to external location as follows:

- 314 ■ The robot can get across, left to right, by traversing the following toggles in the following
 315 order: enter through toggle 1's lower tunnel, down to toggle 5, up to toggle 4's vertical
 316 tunnel, through toggle 1's upper tunnel, around the top to toggle 2's top tunnel, back
 317 down through toggle 4, back out through toggle 5, across through toggle 3's horizontal
 318 tunnel, then through toggle 4's horizontal tunnel, then out through toggle 2's lower
 319 tunnel.
- 320 ■ The robot can get across, right to left, by traversing the following toggles in the following
 321 order: enter through toggle 2's lower tunnel, down to toggle 6, up to toggle 4's vertical
 322 tunnel, through toggle 2's top tunnel, around to toggle 1's top tunnel, down through
 323 toggle 3's vertical tunnel, back out through toggle 6, across through toggle 4's horizontal
 324 tunnel, then through toggle 3's horizontal tunnel, then out through toggle 1's lower
 325 tunnel.
- 326 ■ If the middle toggles are in the opposite orientation, the system is simply mirrored, left
 327 to right, and the traversals are still possible.

328 Next, we wish to establish that the robot cannot cross the middle tunnel in the locked
 329 state. After entering from either middle tunnel location, the only traversable toggles are the
 330 middle toggles. After traversing those, the robot can go no further. The bottom toggles
 331 can't be traversed, so the entire middle region is inaccessible. As a consequence, the opposite
 332 outer toggle's upper pathway can't be accessed. Therefore the robot can only leave via its
 333 original location.

334 We also must establish that if the gadget starts in one of the configurations mentioned,
 335 the robot must leave it in the proper state, and can't leave it in a configuration that wasn't
 336 mentioned. This is straightforward for the bottom tunnel, so we will focus on the middle
 337 two locations.

338 We will show that the accessible configurations of the gadget are exactly as described.
 339 To do so, we will make use of the concept of a cut in a gadget.

340 ► **Lemma 4.5.** *Let A be a connected region of a planar embedding of a gadget system which
 341 does not contain any locations. Then the boundary of A , which we will call a cut, is traversed
 342 an even number of times during any traversal of the construction.*

343 **Proof.** Whenever the boundary of A is crossed, the robot goes from inside A to outside or
 344 vice versa. Since the robot starts a traversal outside A and ends it outside A , it must cross
 345 the boundary an even number of times. ◀

346 The upper pathways of the outer toggles form a cut, and the lower pathways of the outer
 347 toggles form a cut. Thus, the upper pathways of the outer toggles are crossed an even number

348 of times, and the lower pathways are passed an even number of times, so the outer toggles
 349 must be passed an even number of times in total. Thus, the toggles must either be both
 350 oriented in or both out when leaving. However, when leaving the gadget, the outer toggle
 351 which the robot exited through must end up oriented in, so both outer toggles must end up
 352 oriented in.

353 The vertical pathways of the middle toggles form a cut. The horizontal pathways form a
 354 cut. Thus, upon leaving, the middle toggles must have been traversed an even number of
 355 times in total, and hence must end up both left or both right.

356 The upper pathways of the bottom toggles must be passed an even number of times. So
 357 the upper pathways of those toggles must either be both in or both out when leaving the
 358 gadget system.

359 Thus, the gadget system must be left in a state where the outer toggles are oriented in,
 360 the middle toggles are oriented either both left or both right, and the upper pathways of
 361 the bottom toggles are oriented either both in or both out. Therefore, these are exactly the
 362 accessible configurations, as desired.

363 Finally, we show that the robot leaves the gadget in the same state it was entered in,
 364 if it is entered on the middle tunnel. If the robot passes through one of the upper tunnels
 365 of the bottom toggles, when it leaves the region bounded by the bottom toggles' upper
 366 tunnels, it must leave one of the bottom toggle's upper tunnels oriented in. By the parity
 367 constraint, both bottom toggles' upper tunnels will be oriented in, thus leaving the gadget
 368 in the unlocked state. If the central tunnels are entered in the unlocked state, they will be left
 369 in the unlocked state. In the locked state, the upper tunnels of the bottom toggles cannot be
 370 passed, and so the gadget will be left in the locked state.

371 Thus, the construction correctly simulates a NTL. ◀


372 **4.1 2-toggles and non-crossing toggle locks simulate 2-toggle locks**

373 We introduce some new three tunnel objects. There are several
 374 distinct planar topologies of the tunnels in a three tunnel object.
 375 We will focus on the two topologies which can be drawn with
 376 no internal crossing tunnels: three tunnels around the perimeter,
 377 and three tunnels in parallel. We will call the former a "round"
 378 topology, and the latter a "stacked" topology. Note that in the
 379 stacked topology, the order of the tunnels is relevant. In either
 380 topology, if there are multiple toggles, the relative orientation
 381 must still be specified.


382 ▶ **Lemma 4.6.** *2-toggles and noncrossing toggle locks simulate a*
 383 *round antiparallel 2-toggle-lock (RAP2TL) and a round parallel*
 384 *2-toggle-lock (RP2TL).*

385 **Proof.** The construction shown in Figure 7 simulates the behav-
 386 ior of a round antiparallel 2-toggle-lock. It has two externally
 387 accessible states: as shown, and with the middle two gadgets
 388 flipped. These correspond to the 2-toggle of the RAP2TL being
 389 pointed counterclockwise and clockwise respectively.


390 We will demonstrate that this gadget is equivalent to a RAP2TL by examining all possible
 391 traversals. From the two locations that are on the lock tunnel of the NTL, the only possible
 392 traversals are to each other, if the lock tunnel is unlocked. This forms the lock tunnel of the
 393 RAP2TL.



385 **Figure 7** Round antiparallel 2-toggle-lock construction



■ **Figure 8** A round parallel 2-toggle lock is used to construct a stacked antiparallel 2-toggle lock



■ **Figure 9** A noncrossing tripwire lock constructed from an anti-parallel 2-toggle and lock with the lock on the side

394 Traversals from the top left location: The robot must go down and to the right, due to
 395 the orientation of the toggle of the NTL. Then, the robot can go through the C2T, at which
 396 point it is blocked by the orientation of the bottom P2T. Thus, no traversal is possible from
 397 this location in this state.

398 Traversals from the top right location: The robot can go through the C2T, then through
 399 the NTL. At this point, the robot cannot go through the C2T again, because the C2T has
 400 been toggled. Therefore, its only option is to go through the upper P2T and leave at the
 401 top left location. This traversal toggles both of the middle two gadgets, and toggles the
 402 upper P2T twice. Thus, the external state of the gadget is flipped. This is the equivalent of
 403 traversing the upper toggle of the RAP2TL that we are simulating.

404 Traversals from the bottom left location: The robot must go up and to the left, due to
 405 the orientation of the C2T. Then, the robot can go through the NTL. Due to the orientation
 406 of the upper P2T, the robot must now go through the C2T. Now, the robot can leave at the
 407 bottom right location. This traversal toggles both of the middle two gadgets, and toggles
 408 the lower P2T twice. Thus, the external state of the gadget is flipped. This is the equivalent
 409 of traversing the lower toggle of the RAP2TL that we are simulating.

410 Traversals from the bottom right location: The robot is blocked by the orientation of the
 411 C2T. Thus, no traversal is possible from this location in this state.


412 The opposite state is equivalent to a top-bottom mirror reversal, except for a change
 413 in the state of the lock, which does not affect which traversals are possible. Thus, in every
 414 state, this system of gadgets is equivalent to a round antiparallel two-toggle-lock (RAP2TL).

415 Consider the gadget which is the same as the one in Figure 7, except that the bottom P2T
 416 is replaced with a C2T with its toggles allowing traversals from the bottom locations into
 417 the gadget. Clearly, the effect of this change is to swap the roles of the bottom two locations.
 418 As a result, this new construction is a round parallel two-toggle-lock, a RP2TL. ◀


419 ▶ **Lemma 4.7.** *RP2TLs and 2Ts simulate a stacked antiparallel 2-toggle-lock (SAP2TL).*

420 **Proof.** A SAP2TL is a three tunnel gadget where the three tunnels cross the gadget in
 421 parallel, with the two antiparallel toggle tunnels next to each other.

422 Starting with a RP2TL and two C2Ts, we can simulate a SAP2TL as shown in Figure 8.
 423 The lock tunnel is straightforward. The two other traversals are from the top left to the
 424 bottom left, and from the bottom right to the top right. Both of these traversals pass through
 425 every gadget. In the other state, all three gadgets are flipped, and the same traversals are
 426 possible in the opposite direction.



■ **Figure 10** A stacked tripwire-lock-tripwire constructed from non-crossing tripwire locks.



■ **Figure 11** A crossover constructed from stacked tripwire-lock-tripwires

427 Since every state-affecting traversal traverses all gadgets, the states of the three gadgets
 428 always switch together, and the behavior is that of an SAP2TL. Equivalently, by Lemma 2.2,
 429 the system of gadgets is deterministic and reversible, so the three traversals mentioned are
 430 the only ones possible, and the construction simulates a SAP2TL. ◀

431 4.2 2-toggle locks simulate non-crossing wire locks

432 ▶ **Lemma 4.8.** *AP2TLS simulates a NWL.*

433 **Proof.** By connecting the locations of the SAP2TL as shown in Figure 9, we can simulate a
 434 NWL.

435 Each traversal of either connected toggle tunnel flips the state. The connections between
 436 these two tunnels ensure that travel in either direction is always possible. As a result, the
 437 combination of these connected pathways acts as a tripwire, always allowing the robot to
 438 pass in either direction and opening or closing the lock with each traversal. ◀

439 4.3 Non-crossing wire locks simulate crossovers

440 On our way to simulating a crossover, we will simulate another three tunnel gadget, a stacked
 441 tripwire-lock-tripwire (SWLW). Note that the lock tunnel is specifically the center tunnel.

442 ▶ **Lemma 4.9.** *NWLs simulate a stacked tripwire-lock-tripwire (SWLW).*

443 **Proof.** The construction is shown in Figure 10. There are four accessible states of this
 444 gadget, which are any of the states where there is one locked and one unlocked NWL among
 445 the two top NWLs, and one of each among the two bottom NWLs.

446 The states can only be changed by traversing the tripwire tunnels, and doing so flips
 447 both NWLs on the side traversed, maintaining the invariant.

448 If both left NWLs are locked, or both right NWLs are locked, the center tunnel is not
 449 passable. In the other two accessible states, the center tunnel is passable. The two pairs
 450 correspond to the two external states, with the lock locked and unlocked respectively. In any
 451 state, traversing either tripwire moves the gadget to a state with the opposite passability of
 452 the lock tunnel. Thus, this construction simulates a SWLW. ◀

453 ▶ **Lemma 4.10.** *SWLWs simulate a crossover.*

454 **Proof.** The gadget shown in Figure 11 implements a crossover. The robot may always cross
 455 from left to right, right to left, top to bottom and bottom to top, but in no other directions.
 456 There is a single accessible state, the one with all four SWLWs in the unlocked state.

457 When the robot enters from any of the four external locations it has only a single option
 458 up until the point where it reaches the four-way intersection at the center. Upon reaching
 459 this point, the robot has traversed the tripwire tunnels of two of the SWLWs, locking them.
 460 In particular, the SWLWs whose lock tunnels are on the two orthogonal pathways are locked.
 461 For instance, if the robot entered from the top, the left and right pathway's SWLWs would
 462 be locked at this point. As a result, the only way for the robot to continue is to go straight,
 463 passing through the other tripwires of the same two SWLWs, and emerging from the other
 464 side. The robot has completed a crossover traversal, with no other options.

465 Because the robot passed through the tripwires of two SWLWs twice, and only the lock
 466 tunnels of the other two SWLWs, the object is left in its original state, making the state shown
 467 in Figure 11 the only accessible state. This construction correctly simulates a crossover. ◀

468 For the PSPACE-completeness result, we make use of 2-toggle locks and crossovers.
 469 Combining the lemmas in Section 4, we have the result we will make use of:

470 ▶ **Theorem 4.11.** *AP2Ts simulate crossovers and all 2-toggle-locks.*

471 **Proof.** By composing the lemmas in Section 4, we see that AP2Ts simulate crossovers and
 472 RAP2TLs. By using the crossover to effectively rearrange locations, we can simulate an
 473 arbitrary 2-toggle-lock. ◀

474 **5 Everything simulates everything else**

475 The remaining gadgets of interest are each individually (when combined with branching
 476 hallways) sufficient to make motion planning problems PSPACE-complete. Moreover, each
 477 gadget can be simulated by a constant number of each other gadget. To prove this, we give
 478 simple gadgets to show how to construct noncrossing-tripwire-toggles from anti-parallel-2-
 479 toggles, and anti-parallel 2-toggles from each of noncrossing-toggle-locks, noncrossing-wire-
 480 locks, noncrossing-wire-toggles and parallel-2-toggles. We then show that a crossing version
 481 of a gadget can very simply make a non-crossing version of the same gadget.


482 ▶ **Theorem 5.1.** *The 2-toggles, toggle-locks, tripwire-locks and tripwire-toggles, in all orien-*
 483 *tations, can each simulate each other.*

484 **Proof.** We have already established that AP2Ts can simulate P2Ts, C2Ts, NTLs and NWLs
 485 and crossovers. We will establish that:

- 486 ■ AP2Ts can simulate NWTs. Lemma 5.3.
- 487 ■ P2Ts, NTLs, NWTs and NWLs can each simulate AP2Ts. Lemmas 5.4, 5.5, 5.6, 5.7,
 488 respectively.
- 489 ■ C2Ts can simulate P2Ts by Lemma 4.3, and hence AP2Ts as well.
- 490 ■ CTLs can simulate NTLs, CWLs can simulate NWLs, and CWTs can simulate NWTs.
 491 Lemma 5.8.

492 Thus, every gadget can simulate AP2Ts, and AP2Ts can simulate every non-crossing gadget,
 493 as well as crossovers. By combining non-crossing gadgets with crossovers, AP2Ts can simulate
 494 every gadget. This gives a simulation of every gadget by every other gadget, via AP2Ts as
 495 an intermediate step. ◀

496 ▶ **Corollary 5.2.** *Motion planning with any one of the gadgets in Theorem 5.1 (and branching*
 497 *hallways) is PSPACE-complete.*



■ **Figure 12** A noncrossing wire toggle constructed from a toggle, four noncrossing tripwire locks, and two crossovers.

498 **Proof.** Corollary 5.2 follows from Theorem 5.1, which establishes that each gadget can
 499 simulate a AP2T, and Theorem 4.1, which establishes that motion planning with AP2Ts is
 500 PSPACE-complete. ◀

501 ▶ **Lemma 5.3.** *AP2Ts simulate an NWT.*

502 **Proof.** We will construct a NWT as shown in Figure 12. This requires NWLs, crossovers,
 503 and 1-toggles. We already have existing constructions of NWLs and crossovers with AP2Ts.
 504 We can also build a 1-toggle with an AP2T simply by ignoring one of the two tunnels. Thus,
 505 all that's left is to show that the construction successfully simulates a NWT.

506 There are four accessible states: As shown in Figure 12, with all of the NWLs flipped, with
 507 the toggle flipped, and with everything flipped. The first and last correspond to the external
 508 state where the toggle is pointed right, while the other two correspond to the external state
 509 where the toggle is pointed left. The horizontal tunnel corresponds to the toggle, while the
 510 U-shaped tunnel corresponds to the tripwire in the composed gadget. In the state shown in
 511 the figure, the toggle is oriented to the right from the external perspective.


512 Clearly, traversing the U-shaped tunnel will flip all of the tripwires of the NWL, resulting
 513 in a state which corresponds to the opposite external state, as desired.

514 In the state shown in the figure, the horizontal tunnel may be traversed from left to right
 515 along a unique pathway due to the placement of the locks, flipping the toggle along the way.
 516 The orientation of the toggle blocks the right to left traversal. Thus, in this state, the upper
 517 tunnel may be traversed in one direction resulting in an allowed state which corresponds to
 518 the opposite external state, as desired.

519 Placing the toggle in the opposite state is equivalent to a rotation by π of the upper
 520 tunnel, showing this state also correctly simulates an NWT.

521 Flipping the states of all of the NWLs is equivalent to a vertical reflection of the upper
 522 tunnel, showing this state also correctly simulates an NWT. ◀

523 ▶ **Lemma 5.4.** *P2Ts simulate an AP2T.*



■ **Figure 13** Parallel 2-toggles simulate anti-parallel 2-toggles

524 **Proof.** Figure 13 gives a construction of an antiparallel-2-toggle out of parallel-2-toggles.

525 There are two accessible states: As shown, and with the four inner P2Ts flipped. The
 526 former corresponds to the AP2T having a tunnel connecting the left two locations with its
 527 toggle oriented upward, and a tunnel connecting the right locations with its toggle
 528 oriented downward, while the latter corresponds to the two toggles flipped.

529 First, let us examine the bottom right location in the state shown in the figure. After
 530 passing the rightmost P2T, the robot is blocked. No transitions or state changes are possible.
 531 This matches the desired behavior, because the right toggle in the AP2T being simulated is
 532 oriented down.

533 Next, let us examine the top right location in the state shown in Figure 13. After passing
 534 the rightmost P2T, then the upper right P2T, the robot may now either proceed along the
 535 top tunnel, or down to the central loop. In the former case, the robot may pass through the
 536 upper left P2T, but then is blocked. In the later case, the robot may either proceed around
 537 the loop to the left or to the right. If the robot goes to the right, it can pass through the
 538 lower tunnel of the upper right P2T, but then is stuck. If the robot goes to the left, it can
 539 pass through the lower tunnel of the upper left P2T, then the upper tunnel of the lower left
 540 P2T.


541 At this point, the robot may either continue around the loop, or exit the loop downward.
 542 If the robot continues around the loop, it can pass through the upper tunnel of the lower
 543 right P2T, but then is stuck. If it exits the loop, it can either go left or right on the bottom
 544 tunnel. If it goes left, it can pass through the lower tunnel of the lower left P2T, but then is
 545 stuck. If it goes right, it can pass through the lower tunnel of the lower right P2T, then the
 546 lower tunnel of the rightmost P2T, and exit the gadget.

547 Overall, we observe that the robot can make exactly one transition, from top right to
 548 bottom right. The right toggle is traversed twice, and the inner toggles are all traversed
 549 once, leaving the gadget in the other accessible state. No other transition or state change is
 550 possible, from that entrance.


551 Since the gadget is rotationally symmetric about its center, the possible transitions from
 552 the right mirror the possible transitions from the left. Since the other state is simply the state
 553 shown in the figure mirrored top-to-bottom, the transitions described mirror the transitions
 554 in the other state as well. ◀

555 ▶ **Lemma 5.5.** *NTLs simulate an AP2T.*

556 **Proof.** The construction is shown in Figure 14. The two accessible states are the state shown
 557 in the figure and the state with all of the NTLs flipped, but the one-toggles still oriented
 558 inward. These correspond to an AP2T with the top tunnel directed left and bottom tunnel
 559 directed right, and the left-right mirror image.



■ **Figure 14** Noncrossing-toggle-lock simulates anti-parallel-2-toggle



■ **Figure 15** Noncrossing-wire-toggle simulates anti-parallel-2-toggle

560 If the robot enters from the top right, after passing the lock of the top right NTL, it
 561 must pass the upper one-toggle and proceed into the central loop. Since the lower toggle is
 562 directed upward, the robot must eventually leave the central loop via the upper toggle. The
 563 robot may now proceed around the loop. The loop may only be traversed counterclockwise,
 564 and it may only be traversed once. The robot may of course backtrack at any point, but
 565 when it leaves via the upper toggle, it must have either traversed the loop zero or one times.
 566 In the former case, the robot must leave via the top right location, leaving the system in
 567 its original state. In the latter case, the robot must leave via the top left location, as all of
 568 the locks have flipped. Thus, the top tunnel may be traversed via a right to left traversal,
 569 flipping the state, and that is the only traversal in that direction.

570 If the robot enters from the top left, it is immediately blocked by the lock, and no traversal
 571 is possible. Thus, the top tunnel works as desired.

572 Since the gadget possesses rotational symmetry around its center, the bottom tunnel is
 573 exactly the same, allowing only a left to right traversal, flipping the state.


574 The opposite state is the same as the original state except for a left-right right mirror
 575 reversal, so it also functions exactly as desired from the AP2T. ◀

576 ▶ **Lemma 5.6.** *NWTs simulate an AP2T.*

577 **Proof.** A noncrossing wire toggle can simulate an anti-parallel 2-toggle with the simple
 578 construction shown in Figure 15. The direction of each tunnel is dictated by the toggle on
 579 the tunnel, and the wire ensures both toggles are synchronized. Thus when either tunnel is
 580 traversed, both NWTs flip and the direction each tunnel can be traversed flips. ◀

581 ▶ **Lemma 5.7.** *NWLs simulate an AP2T.*

582 **Proof.** The construction of an anti-parallel 2-toggle from non-crossing tripwire locks can
 583 be seen in Figure 16. Note that a 1-toggle can be constructed from an NWL by simply
 584 connecting one location of the wire to one location of the lock. A closed lock will prevent
 585 travel in one direction, but crossing the tripwire in the other direction will open the lock
 586 and allow the robot to proceed. An open lock will allow travel in the other direction. In



■ **Figure 16** Noncrossing-wire-lock simulates anti-parallel-2-toggle


587 the direction starting from the tripwire, the tripwire will close the lock in front of the robot
 588 preventing traversal. In either traversal, the tripwire is crossed, flipping the state.

589 There are two main parts to this gadget, the top and bottom tunnels, and the inner
 590 loop. As with the NTL construction from Lemma 5.5, the 1-toggles ensure that the loop
 591 must be exited from the same place it was entered, which ensures all gadgets on the loop are
 592 traversed the same number of times. Since all wires are on this loop, in a given traversal of
 593 this gadget system, all of the NWLs will change state the same number of times, keeping
 594 them in sync. The upper and lower paths each contain a locked and unlocked tunnel. The
 595 locked portion prevents entry and interaction with the gadget. From the unlocked side, the
 596 robot is able to enter the gadget and flip its state an arbitrary number of times. If the state
 597 is flipped an even number of times, the robot’s only path out is the way it came. If an odd
 598 number of flips have occurred, the robot can now exit through the opposite side of its path,
 599 leaving the gadget in the opposite state.

600 Therefore, the gadget may be traversed right to left along the top tunnel, flipping the state,
 601 and left to right along the bottom tunnel, flipping the state. We have built an AP2T. ◀

602 ▶ **Lemma 5.8.** *CWTs simulate an NWT, CWLs simulate an NWL, CTLs simulate an NTL.*

603 In general, one can very easily simulate a non-crossing
 604 version of a 2-tunnel gadget from the crossing version.
 605 Figure 17 shows a parallel-2-toggle being constructed
 606 from a crossing-2-toggle. The same construction works
 607 for uncrossing the other gadgets we have analyzed, namely
 608 tripwire-toggles, tripwire-locks and toggle-locks. Going
 609 from non-crossing to crossing versions is significantly
 610 more complicated (except in the case of anti-parallel-2-
 611 toggle to crossing-2-toggle) but we are rescued from the
 612 need of such constructions by being able to simulate a general crossover in Lemma 4.10.



■ **Figure 17** Crossing 2-toggles simulate parallel 2-toggle


613 **6 More reasons Zelda is hard**

614 In this section we use this framework to give an alternate proof that The Legend of Zelda:
 615 Oracle of Seasons is PSPACE-complete. Along the way, we will show that motion planning
 616 with reversible deterministic gadgets we call ‘spinners’ is also PSPACE-complete.

617 A k -spinner is a two state deterministic reversible gadget on k locations. In one state,
 618 each location is connected to its neighbor by a directed edge in a clockwise direction. In the
 619 other state, all locations are likewise connected in a counterclockwise direction. A 4-spinner
 620 is shown in Figure 18. The study of 4-spinners was posed by Jeffrey Bosboom due to



■ **Figure 18** Example of a 4-spinner in The Legend of Zelda: Oracle of Seasons.



■ **Figure 19** 4-spinners simulate deterministic forks which simulate crossing 2-toggles

621 their appearance in The Legend of Zelda: Oracle of Seasons. We show that for any $k \geq 4$,
 622 path-planning problems with k -spinners and branching hallways is PSPACE-complete.

623 First, we can take a k spinner and have all but three consecutive locations lead to dead
 624 ends. The remaining three locations form a gadget that we call a deterministic fork. A
 625 deterministic fork is a reversible, deterministic gadget on three locations. In one state, it
 626 allows the robot to go from the center to the right location and return from the left to
 627 the center location. In the other state these directions are reversed. Figure 19 shows the
 628 construction of a crossing 2-toggle from two 4-spinners or equivalently two deterministic
 629 forks.

630 ► **Theorem 6.1.** *For any $k \geq 4$, the path-planning problem with k -spinners and branching*
 631 *hallways is PSPACE-complete.*

632 **Proof.** We construct a deterministic fork by ignoring $k - 3$ of the edges in the spinner.
 633 Two deterministic forks together simulate a crossing 2-toggle as shown in Figure 19. By
 634 Corollary 5.2, the motion planning problem with crossing 2-toggles is PSPACE-complete. ◀

635 ► **Corollary 6.2.** *Determining if a player can beat a level in generalized The Legend of Zelda:*
 636 *Oracle of Seasons is PSPACE-hard.*

637 **Proof.** The Legend of Zelda: Oracle of Seasons contains 4-spinners and requires the player
 638 to navigate from one location to a target location in a grid. Since planar graphs can be laid
 639 out in a grid with only quadratic blowup [2], we can reduce from motion planning problems
 640 with 4-spinners which are PSPACE-complete by Theorem 6.1. ◀

641 The complexity of motion planning with 3-spinners, as well as the two other reversible,
 642 deterministic, 2 state, 3 location gadgets, remains open. Since 2-spinners are the same as an
 643 edge in a graph, this would give a tight characterization for the spinner gadget. The authors
 644 would also be interested to know what other games and puzzles use spinners.

645 **7 General hardness characterization**

646 Here, we tightly characterize the hardness of the motion planning problem with all determin-
 647 istic, reversible, 2-state, k -tunnel gadgets.

648 ► **Theorem 7.1.** *Motion planning with any deterministic, reversible, 2-state, k -tunnel planar*
 649 *gadget (with branching hallways) is PSPACE-complete if and only if the gadget has two toggle*

650 *tunnels, a toggle tunnel and a tripwire tunnel, a toggle tunnel and a lock tunnel or a tripwire*
651 *tunnel and a lock tunnel. Motion planning with all other such gadgets is in P.*

652 First, we provide upper bounds for some classes of simpler gadgets. This shows that, for
653 their category, our hardness results are minimal in the sense that path planning with simpler
654 gadgets in the same class can be solved in P.

655 ► **Theorem 7.2.** *Gadgets with only one state are in NL.*

656 **Proof.** One state gadgets cannot change in any way. Thus they must all be comprised of
657 static descriptions of allowed traversals from one location to another. This can be modeled
658 as a mixed graph. Path planning in mixed graphs is in NL [10]. ◀

659 The only nontrivial gadget on 1 tunnel with two states which is reversible and deterministic
660 is the 1-toggle.

661 ► **Theorem 7.3.** *Motion planning with 1-toggles is in NL.*

662 **Proof.** We reduce this problem to ST connectivity in mixed graphs. To solve this problem
663 we simply treat every 1-toggle as a directed edge pointed in the direction the 1-toggle is
664 initially oriented and then run the standard algorithm. It is obvious that if a solution here
665 exists then a path in the 1-toggle planning problem also exists. What is less clear is that
666 this is sufficient to find any such path.

667 Consider a path which traverses at least one toggle more than once. Consider the last
668 toggle on the path which is traversed more than once. After this toggle is traversed, only
669 toggles which are traversed at most once are on the path. Call this toggle t , and let its final
670 traversal be from u to v . Since t was traversed repeatedly, there was some previous point
671 in the path where the robot was at v , before it traversed t the second-to-last time. Let us
672 create a new path where the robot skips the cycle in the original path from v through t to u ,
673 then eventually back to u through t to v . This path must successfully reach the end, as every
674 toggle after t is traversed at most once, and so is in the same state regardless of whether the
675 cycle is omitted.

676 Thus, under the assumption that there is a path which traverses toggle more than once,
677 there is another, shorter path. Thus, the shortest path must not traverse toggles more than
678 once, and so such a path must exist if any path exists. ◀

679 The remaining two-state two-tunnel deterministic reversible gadgets are also in P. We note
680 that a wire-wire never changes its connectivity and is thus no different than two undirected
681 edges. A lock-lock can never change its state and thus is reducible to a one state gadget,
682 simply zero, one, or two undirected edges. A gadget with a tunnel which does not change
683 and is not changed by the state of the gadget is reducible to two gadgets on one tunnel each,
684 which are in P by Theorem 7.3. This exhausts the 2-state 2-tunnel reversible undirected
685 gadgets.

686 **Proof of Theorem 7.1.** Now, we can characterize all two state, deterministic, reversible
687 gadgets on any number of tunnels.

688 Any gadget with two toggle tunnels, a toggle tunnel and a tripwire tunnel, a toggle tunnel
689 and a lock tunnel or a tripwire and a lock tunnel is sufficient to make motion planning hard,
690 by ignoring all other tunnels and using one of the constructions from this paper.

691 We can divide all other gadgets into three categories: those with tripwires and trivial
692 tunnels, those with locks and trivial tunnels, and those with a single toggle and trivial
693 tunnels. The passability of a tunnel in a gadget with only tripwires and trivial tunnels never

694 changes, making motion planning equivalent to st-connectivity. A gadget with only locks
 695 and trivial tunnels can never have its state change, allowing us to apply Theorem 7.2. A
 696 gadget with a single toggle and some number of trivial tunnels can be treated as a one-toggle
 697 together with some number of undirected edges. Thus, any system of gadgets of these types
 698 is equivalent to a system of 1-toggles and undirected edges. After that, the same argument
 699 as in Theorem 7.3 can be used to solve the motion planning problem in that system. ◀

700 **8** Open Problems / Conclusion

701 This framework for abstract motion planning problems leaves open the question of the
 702 computational complexity of motion planning with many other types of gadgets. One can
 703 examine gadgets with more states, without the tunnel restriction, or without the deterministic
 704 and reversible restrictions. Since this is a vast undertaking with many of the gadgets and
 705 their combinations likely to be uninteresting, we suggest some of the following categories to
 706 be of particular interest.

- 707 ■ 3 spinners are the only size of spinner for which motion planning remains open.
- 708 ■ Three location, 2-state, deterministic, reversible gadgets seem like the obvious ‘simplest’
 709 category of gadgets.
- 710 ■ Are there any sets of purely deterministic and reversible gadgets for which motion planning
 711 is PSPACE-complete (e.g. without branching hallways, which are non-deterministic)?
- 712 ■ What about reversible but nondeterministic gadgets on two tunnels or three locations?

713 There is currently significant partial progress on all of the listed topics. Please contact us
 714 before spending significant time working on the open problems listed to prevent duplication
 715 of effort.

716 Acknowledgments

717 This work grew out of an open problem session from the MIT class on Algorithmic Lower
 718 Bounds: Fun with Hardness Proofs (6.890) from Fall 2014. We particularly thank Jeffrey
 719 Bosboom for posing the problem of analyzing 4-spinners from Legend of Zelda: Oracle of
 720 Seasons (in 2015), for simplifying the 2-state k -tunnels proof, and for other helpful discussions.

721 — References —

- 722 **1** Greg Aloupis, Erik D. Demaine, Alan Guo, and Giovanni Viglietta. Classic Nintendo games
 723 are (computationally) hard. *Theoretical Computer Science*, 586:135–160, 2015. Originally
 724 at FUN 2014.
- 725 **2** H. De Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid.
 726 *Combinatorica*, 10(1):41–51, Mar 1990. URL: <https://doi.org/10.1007/BF02122694>, doi:
 727 10.1007/BF02122694.
- 728 **3** Erik D. Demaine, Martin L. Demaine, Michael Hoffmann, and Joseph O’Rourke. Pushing
 729 blocks is hard. *Computational Geometry: Theory and Applications*, 26(1):21–36, August
 730 2003.
- 731 **4** Erik D. Demaine, Isaac Groszof, and Jayson Lynch. Push-pull block puzzles are hard.
 732 In *Proceedings of the 10th International Conference on Algorithms and Complexity*, pages
 733 177–195, Athens, Greece, May 2017. URL: https://doi.org/10.1007/978-3-319-57586-5_16,
 734 doi:10.1007/978-3-319-57586-5_16.

- 735 **5** Erik D. Demaine, Robert A. Hearn, and Michael Hoffmann. Push-2-f is pspace-complete.
736 In *Proceedings of the 14th Canadian Conference on Computational Geometry*, pages 31–35,
737 Lethbridge, Alberta, Canada, August 12–14 2002.
- 738 **6** Erik D. Demaine, Giovanni Viglietta, and Aaron Williams. Super Mario Bros. is harder/
739 easier than we thought. In *Proceedings of the 8th International Conference on Fun with*
740 *Algorithms*, pages 13:1–13:14, La Maddalena, Italy, June 8–10 2016.
- 741 **7** Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the*
742 *Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.
- 743 **8** Robert A. Hearn and Erik D. Demaine. *Games, Puzzles, and Computation*. A. K. Peters,
744 Ltd., Natick, MA, USA, 2009.
- 745 **9** André G. Pereira, Marcus Ritt, and Luciana S. Buriol. Pull and PushPull
746 are PSPACE-complete. *Theoretical Computer Science*, 628:50–61, 2016. URL:
747 <http://www.sciencedirect.com/science/article/pii/S030439751600205X>, doi:<https://doi.org/10.1016/j.tcs.2016.03.012>.
- 748 **10** Walter J. Savitch. Relationships between nondeterministic and deterministic tape
749 complexities. *Journal of Computer and System Sciences*, 4(2):177–192, 1970.
750 URL: <http://www.sciencedirect.com/science/article/pii/S002200007080006X>, doi:[http://dx.doi.org/10.1016/S0022-0000\(70\)80006-X](http://dx.doi.org/10.1016/S0022-0000(70)80006-X).
- 751 **11** Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th*
752 *Annual ACM Symposium on Theory of Computing*, pages 216–226, San Diego, California,
753 May 1978.
754
755