

1 Distributed controller synthesis for deadlock 2 avoidance

3 **Hugo Gimbert**

4 Université de Bordeaux, CNRS, France

5 **Corto Mascle**

6 Université de Bordeaux, France

7 **Anca Muscholl**

8 Université de Bordeaux, France

9 **Igor Walukiewicz**

10 Université de Bordeaux, CNRS, France

11 — Abstract —

12 We consider the distributed control synthesis problem for systems with locks. The goal is to find
13 local controllers so that the global system does not deadlock. With no restriction this problem is
14 undecidable even for three processes each using a fixed number of locks. We propose two restrictions
15 that make distributed control decidable. The first one is to allow each process to use at most
16 two locks. The problem then becomes Σ_2^P -complete, and even in PTIME under some additional
17 assumptions. The dining philosophers problem satisfies these assumptions. The second restriction is
18 a nested usage of locks. In this case the synthesis problem is NEXPTIME-complete. The drinking
19 philosophers problem falls in this case.

20 **2012 ACM Subject Classification** Theory of computation → Distributed computing models

21 **Keywords and phrases** distributed synthesis, concurrent systems, lock synchronisation, deadlock
22 avoidance

23 **Digital Object Identifier** 10.4230/LIPIcs...1

24 **1** Introduction

25 Synthesis of distributed systems has a big potential since such systems are difficult to write,
26 test, or verify. The state space and the number of different behaviors grow exponentially
27 with the number of processes. This is where distributed synthesis can be more useful than
28 centralized synthesis, because an equivalent, sequential system may be very big. The other
29 important point is that distributed synthesis produces by definition a distributed system,
30 while a synthesized sequential system may not be implementable on a given distributed
31 architecture. Unfortunately, very few settings are known for which distributed synthesis is
32 decidable, and those that we know require at least exponential time.

33 The problem was first formulated by Pnueli and Rosner [28]. Subsequent research showed
34 that, essentially, the only decidable architectures are pipelines, where each process can send
35 messages only to the next process in the pipeline [20, 24, 11]. In addition, the complexity
36 is non-elementary in the size of the pipeline. These negative results motivated the study
37 of distributed synthesis for asynchronous automata, and in particular synthesis with so
38 called causal information. In this setting the problem becomes decidable for co-graph
39 action alphabets [12], and for tree architectures of processes [14, 25]. Yet the complexity
40 is again non-elementary, this time w.r.t. the depth of the tree. Worse, it has been recently
41 established that distributed synthesis with causal information is undecidable for unconstrained
42 architectures [17]. Distributed synthesis for (safe) Petri nets [10] has encountered a similar
43 line of limited advances, and due to [17], is also undecidable in the general case, since it
44 is inter-reducible to distributed synthesis for asynchronous automata [3]. This situation



© Hugo Gimbert, Corto Mascle, Anca Muscholl and Igor Walukiewicz;
licensed under Creative Commons License CC-BY 4.0



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 raised the question if there is any setting for distributed synthesis that covers some standard
 46 examples of distributed systems, and is manageable algorithmically.

47 In this work we consider distributed systems with locks; each process can take or release a
 48 lock from a pool of locks. Locks are one of the most classical concepts in distributed systems.
 49 They are also probably the most frequently used synchronization mechanism in concurrent
 50 programs. We formulate our results in a control setting rather than synthesis – this avoids
 51 the need for a specification formalism. The objective is to find a local strategy for each
 52 process so that the global system does not get stuck. For unrestricted systems with locks we
 53 hit again an undecidability barrier, as for the models discussed above. Yet, we find quite
 54 interesting restrictions making distributed control synthesis for systems with locks decidable,
 55 and even algorithmically manageable.

56 The first restriction we consider is to limit the number of locks available to each process.
 57 The classical example are dining philosophers, where each philosopher has two locks cor-
 58 responding to the left and the right fork. Observe that we do not limit the total number
 59 of processes, or the total number of locks in the system. We show that the complexity of
 60 this synthesis problem is at the second level of the polynomial hierarchy. The problem gets
 61 even simpler when we restrict it to strategies that cannot block a process when all locks are
 62 available. We call them *locally live strategies*. We obtain an NP-algorithm for locally live
 63 strategies, and even a PTIME algorithm when the access to locks is *exclusive*. This means
 64 that once a process tries to acquire a lock it cannot switch to some other action before getting
 65 the lock.

66 The second restriction is nested lock usage. This is a very common restriction in the
 67 literature [19], simply saying that acquiring and releasing locks should follow a stack discipline.
 68 Drinking philosophers [4] are an example of a system of this kind. We show that in this case
 69 distributed synthesis is NEXPTIME-complete, where the exponent in the algorithm depends
 70 only on the number of locks.

71 We formalize the distributed synthesis problem as a control problem [29]. A process is
 72 given as a transition graph where transitions can be local actions, or acquire/release of a
 73 lock. Some transitions are controllable, and some are not. A controller for a process decides
 74 which controllable transitions to allow, depending on the local history. In particular, the
 75 controller of a process does not see the states of other processes. Our techniques are based
 76 on analyzing patterns of taking and releasing locks. In decidable cases there are finite sets of
 77 patterns characterizing potential deadlocks.

78 The notion of patterns resembles locking disciplines [7], a tool frequently used to prevent
 79 deadlocks. An example of a locking discipline is "take the left fork before the right one" in
 80 the dining philosophers problem. Our results allow to check if a given locking discipline may
 81 result in a deadlock, and in some cases even list all deadlock-avoiding locking disciplines.

82 In summary, the main results of this work are:

- 83 ■ Σ_2^P -completeness of the deadlock avoidance control problem for systems where each
 84 process has access to at most 2 locks.
- 85 ■ An NP algorithm when additionally strategies need to be locally live.
- 86 ■ A PTIME algorithm when moreover lock access is exclusive.
- 87 ■ A NEXPTIME algorithm and the matching lower bound for the nested lock usage case.
- 88 ■ Undecidability of the deadlock avoidance control problem for systems with unrestricted
 89 access to locks.

90 Related work

91 Distributed synthesis is an old idea motivated by the Church synthesis problem [5]. Actually,
92 the logic CTL has been proposed with distributive synthesis in mind [6]. Given this long
93 history, there are relatively few results on distributed synthesis. Three main frameworks have
94 been considered: synchronous networks of input/output automata, asynchronous automata,
95 Petri games.

96 The synchronous network model has been proposed by Pnueli and Rosner [27, 28]. They
97 established that controller synthesis is decidable for pipeline architectures and undecidable in
98 general. The undecidability result holds for very simple architectures with only two processes.
99 Subsequent work has shown that in terms of network shape pipelines are essentially the only
100 decidable case [20, 24, 11]. Several ways to circumvent undecidability have been considered.
101 One was to restrict to local specifications, specifying the desired behavior of each automaton
102 in the network separately. Unfortunately, this does not extend the class of decidable
103 architectures substantially [24]. A further-going proposal was to consider only input-output
104 specifications. A characterization, still very restrictive, of decidable architectures for this
105 case is given in [13].

106 The asynchronous (Zielonka automata) model was proposed as a reaction to these negative
107 results [12]. The main hope was that causal memory helps to prevent undecidability arising
108 from partial information, since the synchronization of processes in this model makes them
109 share information. Causal memory indeed allowed to get new decidable cases: co-graph
110 action alphabets [12], connectedly communicating systems [23], and tree architectures [14, 25].
111 There is also a weaker condition covering these three cases [16]. This line of research suffered
112 however from a very recent result showing undecidability in the general case [17].

113 Distributed synthesis in the Petri net model, Petri games, has been proposed recently
114 in [10]. The idea is that some tokens are controlled by the system and some by the environment.
115 Once again causal memory is used. Without restrictions this model is inter-reducible with the
116 asynchronous automata model [3], hence the undecidability result [17] applies. The problem
117 is EXPTIME-complete for one environment token and arbitrary many system tokens [10]. This
118 case stays decidable even for global safety specifications, such as deadlock, but undecidable in
119 general [9]. As a way to circumvent the undecidability, bounded synthesis has been considered
120 in [8, 18], where the bound on the size of the resulting controller is fixed in advance. The
121 approach is implemented in the tool ADAMSYNT [15].

122 The control formulation of the synthesis problem comes from the control theory com-
123 munity [29]. It does not require to talk about a specification formalism, while retaining
124 most useful aspects of the problem. A frequently considered control objective is avoidance
125 of undesirable states. In the distributed context, deadlock avoidance looks like an obvious
126 candidate, since it is one of the most basic desirable properties. The survey [33] discusses the
127 relation between the distributed control problem and Church synthesis. Some distributed
128 versions of the control problem have been considered, also hitting the undecidability barrier
129 very quickly [30, 32, 31, 1].

130 We would like to mention two further results that do not fit into the main threads outlined
131 above. In [34] the authors consider a different synthesis problem for distributed systems: they
132 construct a centralized controller for a scheduler that would guarantee absence of deadlocks.
133 This is a very different approach to deadlock avoidance. Another recent work [2] adds a new
134 dimension to distributed synthesis by considering communication errors in a model with
135 synchronous processes that can exchange their causal memory. The authors show decidability
136 of the synthesis problem for 2 processes.

137 **Outline of the paper**

138 In the next section we define systems with locks, strategies, and the control problem.
 139 We introduce locally live strategies as well as the 2-lock, exclusive, and nested locking
 140 restrictions. This permits to state the main results of the paper. The following three sections
 141 consider systems with the 2-lock restriction. First, we briefly give intuitions behind the
 142 Σ_2^P -completeness in the general case. Section 4 presents an NP algorithm for the distributed
 143 synthesis problem for locally live strategies. Section 5 gives a PTIME algorithm under the
 144 exclusive restriction. Next, we consider the nested locking case, and show that the problem is
 145 NEXPTIME-complete. Finally, we prove that without any restrictions the synthesis problem
 146 for systems with locks is undecidable. Missing proofs are included in the appendix.

147 **2 Main definitions and results**

148 A *lock-sharing system* is a distributed system with components (processes) synchronizing
 149 over locks. Processes do not communicate, but they synchronize using locks from a global
 150 pool. Some transitions of processes are uncontrollable, intuitively the environment decides
 151 if such a transition is taken. The goal is to find a local strategy for each process so that
 152 the entire system never deadlocks. The strategy can observe only local transitions – it does
 153 not see transitions performed by other processes, nor states other processes are in. While
 154 the system is finite state, the challenge comes from the locality of strategies. Indeed, the
 155 unrestricted problem is undecidable. The main contribution of this work are restrictions that
 156 make the problem decidable, and even solvable in PTIME.

157 In this section we define lock-sharing systems, strategies, and the deadlock avoidance
 158 control problem, that is the topic of this paper. We then introduce restrictions on the general
 159 problem and state the main decidability and complexity results.


160 A finite-state *process* p is an automaton $\mathcal{A}_p = (S_p, \Sigma_p, T_p, \delta_p, \text{init}_p)$ with a set of locks T_p
 161 that it can acquire or release. The transition function $\delta_p : S_p \times \Sigma_p \rightarrow \text{Op}(T_p) \times S_p$ associates
 162 with a state from S_p and an action from Σ_p an operation on some lock and a new state; it is
 163 a partial function. The lock operations are acquire (acq_t) or release (rel_t) some lock t from
 164 T_p , or do nothing: $\text{Op}(T_p) = \{\text{acq}_t, \text{rel}_t \mid t \in T_p\} \cup \{\text{nop}\}$. Figure 1 gives an example.

165 A *local configuration* of process p is a state from S_p together with the locks p currently
 166 owns: $(s, B) \in S_p \times 2^{T_p}$. The initial configuration of p is $(\text{init}_p, \emptyset)$, namely the initial
 167 state with no locks. A transition between configurations $(s, B) \xrightarrow{a, \text{op}} (s', B')$ exists when
 168 $\delta_p(s, a) = (\text{op}, s')$ and one of the following holds:

- 169 ■ $\text{op} = \text{nop}$ and $B = B'$;
- 170 ■ $\text{op} = \text{acq}_t$, $t \notin B$ and $B' = B \cup \{t\}$;
- 171 ■ $\text{op} = \text{rel}_t$, $t \in B$, and $B' = B \setminus \{t\}$.

172 A *local run* $(a_1, \text{op}_1)(a_2, \text{op}_2) \cdots$ of \mathcal{A}_p is a finite or infinite sequence over $\Sigma_p \times \text{Op}(T_p)$ such that
 173 there exists a sequence of configurations $(\text{init}_p, \emptyset) = (s_0, B_0) \xrightarrow{(a_1, \text{op}_1)}_p (s_1, B_1) \xrightarrow{(a_2, \text{op}_2)}_p \cdots$
 174 While the run is determined by the sequence of actions, we prefer to make lock operations
 175 explicit. We write Runs_p for the set of runs of \mathcal{A}_p .

176 A *lock-sharing system* $\mathcal{S} = ((\mathcal{A}_p)_{p \in \text{Proc}}, \Sigma^s, \Sigma^e, T)$ is a set of processes together with a
 177 partition of actions between *controllable* and *uncontrollable* actions, and a set T of locks. We
 178 have $T = \bigcup_{p \in \text{Proc}} T_p$, for the set of all locks. Controllable and uncontrollable actions belong
 179 to the system and to the environment, respectively. We write $\Sigma = \bigcup_{p \in \text{Proc}} \Sigma_p$ for the set of
 180 actions of all processes and require that (Σ^s, Σ^e) partitions Σ . The sets of states and action
 181 alphabets of processes should be disjoint: $S_p \cap S_q = \emptyset$ and $\Sigma_p \cap \Sigma_q = \emptyset$ for $p \neq q$. The sets
 182 of locks are not disjoint, in general, since processes may share locks.



■ **Figure 1** A dining philosopher p . Dashed transitions are controllable.

183 ▶ **Example 1.** The dining philosophers problem can be formulated as control problem
 184 for a lock-sharing system $\mathcal{S} = ((\mathcal{A}_p)_{p \in Proc}, \Sigma^s, \Sigma^e, T)$. We set $Proc = \{1, \dots, n\}$ and
 185 $T = \{t_1, \dots, t_n\}$ as the set of locks. For every process $p \in Proc$, process \mathcal{A}_p is as in Figure 1,
 186 with the convention that $t_{n+1} = t_1$. Actions in Σ^s are marked by dashed arrows. These
 187 are controllable actions. The remaining actions are in Σ^e . Once the environment makes a
 188 philosopher p hungry, she has to get both the left (t_p) and the right (t_{p+1}) fork to eat. She
 189 may however choose the order in which she takes them; actions *left* and *right* are controllable.
 190

191 A *global configuration* of \mathcal{S} is a tuple of local configurations $C = (s_p, B_p)_{p \in Proc}$ provided
 192 the sets B_p are pairwise disjoint: $B_p \cap B_q = \emptyset$ for $p \neq q$. This is because a lock can be
 193 taken by at most one process at a time. The initial configuration is the tuple of initial
 194 configurations of all processes.

195 Such systems are *asynchronous*, with transitions between two configurations done by a
 196 single process: $C \xrightarrow{(p, a, op)} C'$ if $(s_p, B_p) \xrightarrow{(a, op)}_p (s'_p, B'_p)$ and $(s_q, B_q) = (s'_q, B'_q)$ for every
 197 $q \neq p$. A global run is a sequence of transitions between global configurations. Since our
 198 systems are deterministic we usually identify a global run by the sequence of transition labels.
 199 A global run w *determines a local run* of each process: $w|_p$ is the subsequence of p 's actions
 200 in w .

201 A *control strategy* for a lock-sharing system is a tuple of local strategies, one for each
 202 process: $\sigma = (\sigma_p)_{p \in Proc}$. A *local strategy* σ_p says which actions p can take depending on a
 203 local run so far: $\sigma_p : Runs_p \rightarrow 2^{\Sigma_p}$, provided $\Sigma^e \cap \Sigma_p \subseteq \sigma_p(u)$, for every $u \in Runs_p$. This
 204 requirement says that a strategy cannot block environment actions.

205 A local run u of a system *respects* σ_p if for every non-empty prefix $v(a, op)$ of u , we have
 206 $a \in \sigma_p(v)$. Observe that local runs are affected only by the local strategy. A global run w
 207 respects σ if for every process p , the local run $w|_p$ respects σ_p . We often say just σ -run,
 208 instead of "run respecting σ ".

209 As an example consider the system for two philosophers from Example 1. Suppose that
 210 both local strategies always say to take the *left* transition. So $hungry^1, left^1, acq_{t_1}^1, acq_{t_2}^1$
 211 is a local run of process 1 respecting the strategy; similarly $hungry^2, left^2, acq_{t_2}^2, acq_{t_1}^2$ for
 212 process 2. (We use superscripts to indicate the process doing an action.) The global
 213 run $hungry^1, hungry^2, left^1, left^2, acq_{t_1}^1, acq_{t_2}^2$ respects the strategy and blocks, since each
 214 philosopher needs a lock the other one owns.

215 ▶ **Definition 2** (Deadlock avoidance control problem). A σ -run w leads to a deadlock in σ if
 216 w cannot be prolonged to a σ -run. A control strategy σ is *winning* if no σ -run leads to a
 217 deadlock in σ . The deadlock avoidance control problem is to decide if for a given system
 218 there is some winning control strategy.

219 In this work we consider several variants of the deadlock avoidance control problem.
 220 Maybe surprisingly, in order to get more efficient algorithms we need to exclude strategies
 221 that can block a process by itself:

222 ► **Definition 3** (Locally live strategy). *A local strategy σ_p for process p is locally live if every
 223 local σ_p -run u can be prolonged to a σ_p -run: there is some $b \in \Sigma_p$ such that ub is a local run
 224 respecting σ_p . A strategy σ is locally live if every local strategy is so.*

225 In other words, a locally live strategy guarantees that a process does not block if it runs
 226 alone. Coming back to Example 1: a strategy always offering one of the *left* or *right* actions
 227 is locally live. A strategy that offers none of the two is not. Observe that blocking one
 228 process after the hungry action is a very efficient strategy to avoid a deadlock, but it is not
 229 the intended one. This is why we consider locally live to be a desirable property rather than
 230 a restriction.

231 Note that being locally live is not exactly equivalent to a strategy always proposing at
 232 least one outgoing transition. In our semantics, a process blocks if it tries to acquire a lock
 233 that it already owns, or to release a lock it does not own. But it becomes equivalent thanks
 234 to the following remark:

235 ► **Remark 4.** We can assume that each process keeps track in its state which locks it owns.
 236 Note that this assumption does not compromise the complexity results when the number of
 237 locks a process can access is fixed. We will not use this assumption in Section 6, where a
 238 process can access arbitrarily many locks (in nested fashion).

239 Without any restrictions our synthesis problem is undecidable.

240 ► **Theorem 5.** *The deadlock avoidance control problem for lock-sharing systems is undecidable.
 241 It remains so when restricted to locally live strategies.*

242 We propose two cases when the control problem becomes decidable. The two are defined
 243 by restricting the usage of locks.

244 ► **Definition 6** (2LSS). *A process $\mathcal{A}_p = (S_p, \Sigma_p, T_p, \delta_p, init_p)$ uses two locks if $|T_p| = 2$. A
 245 system $\mathcal{S} = ((\mathcal{A}_p)_{p \in Proc}, \Sigma^s, \Sigma^e, T)$ is 2LSS if every process uses two locks.*

246 Note that in the above definition we do not bound the total number of locks in the system,
 247 just the number of locks per process. The process from Figure 1 is 2LSS. Our first main
 248 result says that the control problem is decidable for 2LSS.

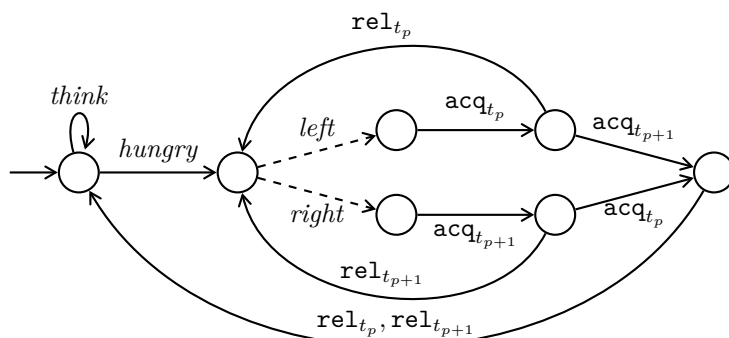
249 ► **Theorem 7.** *The deadlock avoidance control problem for 2LSS is Σ_2^P -complete.*

250 For the lower bound we use strategies that take a lock and then block. This does not
 251 look like a very desired behavior, and this is the reason for introducing the concept of locally
 252 live strategies. The second main result says that restricting to locally live strategies helps.

253 ► **Theorem 8.** *The deadlock avoidance control problem for 2LSS is in NP when strategies
 254 are required to be locally live.*

255 We do not know if the above problem is in PTIME. We can get a PTIME algorithm under
 256 one more assumption.

257 ► **Definition 9** (Exclusive systems). *A process p is exclusive if for every state $s \in S_p$: if s
 258 has an outgoing transition with some acq_t operation then all outgoing transitions of s have
 259 the same acq_t operation. A system is exclusive if all its processes are.*



■ **Figure 2** A flexible philosopher p . She can release a fork if the other fork is not available.

260 ▶ **Example 10.** The process from Figure 1 is exclusive, while the one from Figure 2 is not.
 261 The latter has a state with one $\text{acq}_{t_{p+1}}$ and one rel_{t_p} outgoing transition. Observe that in
 262 this state the process cannot block, and has the possibility to take a lock at the same time.
 263 Exclusive systems do not have such a possibility, so their analysis is much easier.

264 ▶ **Theorem 11.** *The deadlock avoidance control problem for exclusive 2LSS is in PTIME,*
 265 *when strategies are required to be locally live.*

266 Without local liveness, the problem stays Σ_2^P -hard for exclusive 2LSS. Our last result uses a
 267 classical restriction on the usage of locks:

268 ▶ **Definition 12** (Nested-locking). *A local run is nested-locking if the order of acquiring and*
 269 *releasing locks in the run respects a stack discipline, i.e., the only lock a process can release*
 270 *is the last one it acquired. A local strategy is nested-locking if all local runs respecting the*
 271 *strategy are nested-locking. A strategy is nested-locking if all local strategies are nested-locking.*

272 The process from Figure 1 is nested-locking, while the one from Figure 2 is not.

273 ▶ **Theorem 13.** *The deadlock avoidance control problem is NEXPTIME-complete when*
 274 *strategies are required to be nested-locking.*

275 3 Two locks per process

276 We give some intuitions as to why the deadlock avoidance problem for 2LSS is Σ_2^P -complete
 277 (Theorem 7), the details can be found in Appendix A.

278 When every process uses only two locks there are only few patterns of local lock usage
 279 that are relevant for deadlocks. A finite local run u of process p using locks t_1, t_2 can be of
 280 one of the following four types:

- 281 ■ p owns both locks at the end of u ;
- 282 ■ p owns no lock at the end of u ;
- 283 ■ p owns only one lock, say t_1 , at the end of u , and the last lock operation of u is acq_{t_1} ;
- 284 ■ p owns only one lock, say t_1 , at the end of u , and the last lock operation of u is rel_{t_2} .

285 A *pattern* of a run is its type, and the set of available actions at the end. If a run reaches a
 286 deadlock then the only available actions are to acquire locks owned by other processes.

287 We fix a 2LSS $(\{\mathcal{A}_p\}_{p \in \text{Proc}}, \Sigma^s, \Sigma^e, T)$ over the set of processes Proc . We assume that it
 288 satisfies Remark 4.

289 Given a strategy $\sigma = (\sigma_p)_{p \in Proc}$, we call a local σ -run *risky* if it ends in a state from
 290 which every outgoing action allowed by σ acquires some lock (this includes states with no
 291 outgoing transition). A local σ -run is *neutral* if it ends in a configuration (s, B) with $B = \emptyset$.

292 ► **Definition 14.** We define the pattern of a *risky* σ_p -run u_p as follows. Let T_{owns} be the set
 293 of locks that p owns after executing u_p and T_{blocks} the set of locks that outgoing transitions
 294 allowed by σ_p after u_p need to acquire.

295 The pattern of u_p is the tuple $(T_{owns}, T_{blocks}, ord)$ with:

- 296 ■ If u_p is of the form $u_1(a, \mathbf{acq}_{t_1})u_2(b, \mathbf{rel}_{t_2})u_3$ with no action on t_1 in u_2 and no action
 297 on either t_1 or t_2 in u_3 then $ord = (t_1, t_2)$.
- 298 ■ Otherwise $ord = \perp$.

299 Note that in light of Remark 4, T_{owns} and T_{blocks} are necessarily disjoint. Furthermore
 300 if ord is of the form (t_1, t_2) then $T_{owns} = \{t_1\}$, and either $T_{blocks} = \emptyset$ or $T_{blocks} = \{t_2\}$.

301 A strategy $\sigma = (\sigma_p)_{p \in Proc}$ respects a family of sets of patterns $(\mathit{Patt}_p)_{p \in Proc}$ if for all
 302 $p \in Proc$, the patterns of all risky σ_p -runs belong to Patt_p .

303 In this definition, T_{owns} and T_{blocks} serve as witnesses of deadlock configurations, in which
 304 all required locks are owned by another process, and no lock is owned by two different processes.
 305 Further, the ord component indicates the fourth case described before the definition.

306 Our key result in this part is Lemma 15. It gives simple, necessary and sufficient,
 307 conditions on the family of patterns of local σ -runs $(u_p)_{p \in Proc}$ that lead to a deadlock under
 308 a suitable scheduling. The difficulty is to verify if there exists a global run which is a
 309 combination of those local runs. For that, all processes must own disjoint sets of locks at the
 310 end. The rest can be inferred from the types of runs listed above.

311 We describe how to schedule local runs into a global one depending on the four types
 312 listed before Definition 14.

- 313 ■ In the first case we can assume that p 's run is scheduled at the end of the global run, as
 314 it ends up keeping both locks anyway, so no other process will use them after p .
- 315 ■ In the second case, we can assume that p 's run is scheduled at the beginning of the global
 316 run, as it is neutral.
- 317 ■ In the third case, we can split p 's run in two parts: a first, neutral part which can be
 318 scheduled at the beginning, and a second part in which p acquires t_1 and there is no lock
 319 operation afterwards. The second part can be scheduled at the end, because no other
 320 process will use t_1 after p .
- 321 ■ In the final case, p acquires t_1 , never releases it but later uses t_2 . This can be a problem
 322 if for instance another process does the same with t_1 and t_2 reversed. The first process
 323 that takes its first lock would prevent the other from finishing its local run. We express
 324 these constraints by requiring the existence of a global order in which process take locks
 325 without releasing them.

326 ► **Lemma 15.** Let $\sigma = (\sigma_p)_{p \in Proc}$ be a control strategy. For all p let Patt_p be the set of
 327 patterns of local risky σ_p -runs of p . The control strategy σ is **not** winning if and only if there
 328 exists for each p a pattern $(T_{owns}^p, T_{blocks}^p, ord_p) \in \mathit{Patt}_p$ such that:

- 329 ■ $\bigcup_{p \in Proc} T_{blocks}^p \subseteq \bigcup_{p \in Proc} T_{owns}^p$,
- 330 ■ the sets T_{owns}^p are pairwise disjoint,
- 331 ■ there exists a total order \leq on T such that for all p , if $ord_p = (t, t')$ then $t \leq t'$.

332 **Proof.** Suppose σ is not winning, let u be a run ending in a deadlock. For each process p let
 333 u_p be the corresponding local run. The local run u_p is risky, as otherwise u_p could be extended
 334 in a longer run consistent with σ . Thus u_p has a pattern $(T_{owns}^p, T_{blocks}^p, ord_p) \in Patt_p$.

335 We check that those patterns $(u_p)_{p \in Proc}$ meet the requirements of the lemma. Clearly as
 336 we are in a deadlock, all locks that some process wants are taken, hence the first condition
 337 is satisfied. Furthermore, no two processes can own the same lock, implying the second
 338 condition. Finally, let \leq be a total order on locks given by the order of the last operations
 339 on each lock in u : we set $t \leq t'$ iff the last operation on t in u is before the last one on t' . Let
 340 p be a process, and suppose ord_p is (t, t') . Then u_p has the form $u_1(a, \mathbf{acq}_t)u_2(b, \mathbf{rel}_{t'})u_3$
 341 with no action on t in u_2 or u_3 . Hence, $t \leq t'$.

342 The other direction is a bit more complicated. Suppose that for each p there is a pattern
 343 $(T_{owns}^p, T_{blocks}^p, ord_p) \in Patt_p$ such that those patterns satisfy all three conditions of the
 344 lemma. Let \leq be a suitable total order on locks for the third condition, and let $<$ be its
 345 strict part. For every p there exists a risky local run u_p yielding the chosen pattern for p .

346 We start by executing all neutral runs u_p one by one in some order. All locks are free
 347 after these executions.

348 For all p such that $T_{owns}^p = \{t\}$ and $ord_p = \perp$, we can decompose u_p as $u_1(a, \mathbf{acq}_t)u_2$
 349 with no action on locks in u_2 . We execute all runs u_1 , which are neutral and thus leave all
 350 locks free after execution.

351 Finally, we execute all u_p such that $ord_p \neq \perp$ in increasing order on the first component
 352 of ord_p according to \leq . For all such p , let $(t, t') = ord_p$, so we have $T_{owns}^p = \{t\}$ and $t < t'$.
 353 As all T_{owns}^p are disjoint, before executing u_p all locks greater or equal to t according to \leq
 354 are free. In particular, t and t' are free, thus we can execute u_p . In the end all locks are free
 355 except the ones belonging to T_{owns}^p for those processes p .

356 Now we execute the remaining part of the u_p with $T_{owns}^p = \{t\}$ and $ord_p = \perp$ (referred
 357 to as $(a, \mathbf{acq}_t)u_2$ before). Those runs do not contain any action on locks besides the first
 358 acquire. As all T_{owns}^p are disjoint, the locks they acquire are free, hence all those runs can
 359 be executed.

360 The remaining runs are the ones such that $T_{owns}^p = \{t, t'\}$. As all T_{owns}^p are disjoint,
 361 both these locks are free, hence u_p can be executed as p can only use these two locks.

362 We have combined all local runs into one global run reaching a configuration where
 363 all processes have to acquire a lock from $\bigcup_{p \in Proc} T_{blocks}^p$ to keep running, and all locks in
 364 $\bigcup_{p \in Proc} T_{owns}^p$ are taken. As $\bigcup_{p \in Proc} T_{blocks}^p \subseteq \bigcup_{p \in Proc} T_{owns}^p$, we have reached a deadlock.
 365 ◀

366 The algorithm for Theorem 7 proceeds in four phases:

- 367 ■ guess a set of patterns $Patt_p$, one for each process p ,
- 368 ■ check that there are local strategies σ_p such that the patterns of all runs belong to $Patt_p$,
- 369 ■ let the adversary guess a pattern in each $Patt_p$,
- 370 ■ check whether those patterns satisfy the conditions of Lemma 15.

371 The alternation between guessing and adversarial guessing yields a Σ_2^P algorithm.

372 The lower bound is obtained by a reduction from $\exists\forall$ -SAT. The system controls existential
 373 variables, the environment controls universal ones. There are two locks for each variable,
 374 acquiring one of them is interpreted as choosing the value of the variable. The processes
 375 enforcing the choice are displayed in Figure 3 in the appendix. Note that this construction
 376 relies on processes that take a lock and then block on their own in states with no outgoing
 377 transitions. In the following section we will forbid such unnatural behavior by considering
 378 only locally live strategies.

379 We use some extra processes to enforce that the system wins if and only if the valuation
 380 given by the choices of the two players satisfies the SAT formula. The interesting part is
 381 that even though it looks like the guessing values of variables is done concurrently by the
 382 system and the environment, the whole setting enforces a $\exists\forall$ dependency.

383 **4 Two locks per process with locally live strategies**

384 We describe how to solve the control problem for 2LSS and locally live strategies in NP, as
 385 stated in Theorem 8. The full proof is in Appendix B.

386 We fix a 2LSS satisfying the assumption discussed in Remark 4. We will show that the
 387 relevant information about a strategy σ can be formalized as a finite lock graph G_σ and a
 388 lockset family $Locks_\sigma$; the latter is a family of sets of sets of locks (see definitions below).
 389 This information is very similar to the one described by patterns in the previous section.
 390 As we work with locally live strategies, the set of possible patterns of local runs is more
 391 restricted and we can view this more conveniently as a graph.

392 Our algorithm first guesses an abstract lock graph G and lockset family $Locks$. Then it
 393 performs two checks:

394 **Step 1** check if there is some strategy σ with $G = G_\sigma$ and $Locks = Locks_\sigma$, and

395 **Step 2** check if there is no deadlock scheme for G and $Locks$ (see Definition 21 below).

396 A deadlock scheme is some kind of forbidden situation. It is easy to get a co-NP algorithm for
 397 the second step: just guess the scheme and check that it has the right shape. The challenge
 398 is to do this in PTIME. This is necessary if we want to get an NP algorithm.

399 We introduce now some notions in order to define G_σ and $Locks_\sigma$ conveniently. Consider
 400 a local run u of a process p :

$$401 \quad (init_p, \emptyset) = (s_0, B_0) \xrightarrow{(a_1, op_1)}_p (s_1, B_1) \cdots \xrightarrow{(a_i, op_i)}_p (s_i, B_i) .$$

402 We say that u has set of locks B if $B = B_i$. A σ_p -run u is B -locked by the local strategy σ_p
 403 if every transition in $\sigma_p(u)$ has as operation acq_t for some $t \in B$. Process p is B -lockable by
 404 σ_p if it has a neutral, B -locked σ_p -run.

405 The intuition is that in order to get a deadlock, a B -lockable process can be scheduled
 406 first. It can do a run leading to a state where it requires some of the locks in B without
 407 holding any locks. So, the process will be blocked if we ensure that all locks in B are already
 408 taken. For example, consider the process in Figure 1. The run *hungry, left* is $\{t_p\}$ -locked,
 409 as the unique next action is acq_{t_p} . The process is $\{t_p\}$ -lockable by σ_p if e.g. σ_p always
 410 chooses the *left* action. Indeed, in this case the run *hungry, left* is a neutral σ_p -run, which
 411 is $\{t_p\}$ -locked. Process p is not $\{t_{p+1}\}$ -lockable by a strategy σ_p choosing always the *left*
 412 action, as there is no neutral σ_p -run leading to $acq_{t_{p+1}}$.

413 **► Definition 16** (Lockset family $Locks_\sigma$). A lockset for a local strategy σ_p is a set $L_p \subseteq 2^{T_p}$
 414 of sets B such that p is B -lockable by σ_p . A lockset family for σ is $Locks_\sigma = (L_p)_{p \in Proc}$.

415 **► Definition 17** (Lock graph G_σ). For a strategy σ , a lock graph $G_\sigma = \langle T, E_\sigma \rangle$ has an edge
 416 $t_1 \xrightarrow{p} t_2$ whenever there is some σ_p -run u of p that has $\{t_1\}$ and is $\{t_2\}$ -locked. If there is
 417 such a run u where the last lock operation in u is acq_{t_1} then the edge is called green, and
 418 otherwise it is called blue.

419 We will say that σ allows a blue edge $t_1 \xrightarrow{p} t_2$ or a green edge $t_1 \xrightarrow{p} t_2$. We write $t_1 \xrightarrow{p} t_2$
 420 when the color of the edge is irrelevant.

421 For example, a strategy choosing the *left* action in Figure 1 yields the green edge $t_p \xrightarrow{p} t_{p+1}$.
 422 Lockset families say on which sets of locks each process can block while not holding any
 423 lock. An edge $t_1 \xrightarrow{p} t_2$ in the lock graph corresponds to a run of p where P owns lock t_1 (the
 424 source of the edge) and waits for the other lock t_2 (the target of the edge).

425 A lockset represents a run of the second type in the previous section, a green edge a run
 426 of the third type, and a blue edge a run of the fourth type with no similar run of the third
 427 type. The first type cannot appear in a deadlock when strategies are locally live, as processes
 428 always have an available action.

429 Since we have assumed nothing about how strategies are given, it is not clear how to
 430 compute G_σ . Instead of restricting to, say, finite memory strategies, we will work with
 431 arbitrary lock graphs and lockset families. This is possible thanks to Lemma 19 below, that
 432 allows to check if a graph is the lock graph of some strategy. For this we need to define
 433 lockset families and lock graphs abstractly. Notice that the size of both these objects is
 434 bounded, as the set of locks per process is fixed for 2LSS.

435 ► **Definition 18.** A lockset family is a tuple of sets of locks indexed by processes $(L_p)_{p \in Proc}$,
 436 with $L_p \subseteq 2^{T_p}$. A lock graph is an edge-labeled graph $G = \langle T, E \subseteq T \times Proc \times \{blue, green\} \times T \rangle$
 437 where nodes are locks from the set T and every edge is labeled by a process and a color. A
 438 cycle in G is called proper if all its edges are labeled by different processes. It is denoted as
 439 green if it contains at least **one** green edge; otherwise, so if **all** edges are blue, it is denoted
 440 blue.

441 At this point we have enough notions to carry out the first step on page 10.

442 ► **Lemma 19.** Given a lock graph G and a lockset family $Locks$, it is decidable in PTIME if
 443 there is a locally live strategy σ such that $G = G_\sigma$ and $Locks_\sigma = Locks$.

444 The proof is by reduction to model-checking a fixed-size MSOL formula over a given
 445 regular tree. For every process p we need to check if there is a local strategy σ_p satisfying
 446 the conditions imposed by G and $Locks = (L_p)_{p \in Proc}$. Consider the regular tree of all local
 447 runs of process p . The formula says that there is a strategy tree inside this regular tree such
 448 that L_p contains exactly those sets B such that the subtree has some neutral, B -locked path;
 449 and for every edge in G labelled by p there is a path of the required shape in the subtree.
 450 This can be expressed by an MSOL formula of constant size, as the process uses only 2 locks.
 451 From the MSOL formula we get a tree automaton of constant size. The emptiness check of
 452 its product with the tree automaton accepting the unfolding of the automaton \mathcal{A}_p can be
 453 done in PTIME.

454 In the rest of the section we discuss the second step. We first define a Z -deadlock scheme
 455 for some set Z of locks. Intuitively, this is a situation showing that there is a run blocking
 456 all locks in Z . Then a deadlock scheme is a Z -deadlock scheme for some Z big enough to
 457 block all processes.

458 ► **Definition 20** (Z -deadlock scheme). Let $G = \langle T, E \rangle$ be a lock graph, $Locks = (L_p)_{p \in Proc}$
 459 a lockset family, and Z a set of locks. We define $Proc_Z$ as the set of processes whose both
 460 accessible locks are in Z , $Proc_Z = \{p \in Proc : T_p \subseteq Z\}$. A Z -deadlock scheme is a function
 461 $ds_Z : Proc_Z \rightarrow E \cup \{\perp\}$ such that:

- 462 ■ For all $p \in Proc_Z$, if $ds_Z(p) \neq \perp$ then $ds_Z(p)$ is an edge of G labeled by p .
- 463 ■ If $p \in Proc_Z$ and $L_p = \emptyset$ then $ds_Z(p) \neq \perp$.
- 464 ■ For all $t \in Z$ there exists a unique $p \in Proc_Z$ such that $ds_Z(p)$ is an outgoing edge from t .
- 465 ■ The subgraph of G , restricted to $ds_Z(Proc_Z)$ does not contain any blue cycle.

466 The main point of this definition is that for every lock in Z there is an outgoing edge in
 467 ds_Z . Intuitively, it means that we have a run where every lock from Z is taken, and every
 468 process in $Proc_Z$ requires a lock from Z .

469 ► **Definition 21** (Deadlock scheme). *A deadlock scheme for G and $Locks = (L_p)_{p \in Proc}$ is a*
 470 *Z -deadlock scheme such that for every process $p \in Proc \setminus Proc_Z$ there is $B \in L_p$ with $B \subseteq Z$.*

471 Thus a deadlock scheme represents a situation where all processes are blocked, since every
 472 process not in $Proc_Z$ can be brought into a state where it needs a lock from Z , but all these
 473 locks are taken.

474 The next lemma says that the absence of deadlock schemes characterizes winning strategies.
 475 We could reuse the patterns defined above to obtain a shorter proof but we prefer to give a
 476 slightly longer but elementary one.

477 ► **Lemma 22.** *A locally live control strategy σ is winning if and only if there is no deadlock*
 478 *scheme for its lock graph G_σ and its lockset family $Locks_\sigma$.*

479 **Proof.** Suppose σ is not winning. Then there exists a global σ -run u leading to a deadlock.
 480 As a consequence, in the deadlock configuration all processes must be trying to acquire some
 481 lock that is already taken.

482 We then construct a deadlock scheme (BT, ds) as follows. Let BT be the set of locks
 483 taken in the deadlock configuration, and for all $p \in Proc$, define $ds(p)$ as:

- 484 ■ \perp if p does not own any lock in the deadlock configuration,
- 485 ■ $t_1 \xrightarrow{p} t_2$ if p owns t_1 and is trying to acquire t_2 in the deadlock configuration (the color
 486 of the edge is determined by the run, it is irrelevant for the argument).

487 Clearly for all $p \in Proc$ the value $ds(p)$ is either \perp or a p -labeled edge of the lock graph
 488 G_σ .

489 Suppose $ds(p) = \perp$, and let t_1, t_2 be the two locks accessible by p . As the final configuration
 490 is a deadlock, all actions allowed by σ_p are necessarily acq_{t_1} or acq_{t_2} . So p is $\{t_1, t_2\}$ -lockable.
 491 Furthermore, as we are in a deadlock, the lock(s) blocking p are in BT (if they were free, p
 492 would be able to advance), therefore p is BT -lockable.

493 For every $t \in BT$, there is a process p holding t in the final configuration. As we are in
 494 a deadlock, p is trying to acquire its other accessible lock t' (recall that the definition of
 495 control strategy demands that at least one action be available to each process at all times).
 496 Thus $ds(p)$ is an edge from t to t' . Furthermore t' cannot be free as we are in a deadlock,
 497 thus $t' \in BT$. There are no other outgoing edges from t as no other process can hold t while
 498 p does.

499 Finally let $t_1 \xrightarrow{p_1} t_2 \cdots \xrightarrow{p_k} t_{k+1}$ be a cycle with $t_1 = t_{k+1}$ in the subgraph of G_σ restricted
 500 to BT and $ds(Proc)$. One of the locks t_i was the last lock taken in the run u (say by process
 501 p_i). We show now by contradiction that the edge $t_i \xrightarrow{p_i} t_{i+1}$ is green. If p_i would have
 502 released t_{i+1} after the last acq_{t_i} in u , then p_{i+1} would have done its last $\text{acq}_{t_{i+1}}$ later, a
 503 contradiction. The subgraph of G_σ restricted to BT and $ds(Proc)$ has therefore no blue
 504 cycles, therefore (BT, ds) is a deadlock scheme.

505 For the other direction, suppose we have a deadlock scheme (BT, ds) for the lock graph
 506 G_σ . As $(BT, ds(Proc))$ does not contain a blue cycle, we can pick a total order \leq on locks
 507 such that for all blue edges $t_1 \xrightarrow{p} t_2 \in ds(Proc)$, we have $t_1 \leq t_2$.

508 By definition of the lock graph, for each process $p \in Proc$ we can take a local run u_p of
 509 \mathcal{A}_p respecting σ with the following properties.

- 510 ■ If $ds(p) = \perp$ then p is BT -lockable. So there exists a neutral run u_p leading to a state
 511 where all outgoing transitions require locks from BT .

512 ■ If $ds(p) = t_p^1 \xrightarrow{p} t_p^2$ then there is u_p of the form $u_p^1(a, \text{acq}_{t_p^1})u_p^2(a', \text{acq}_{t_p^2})$ without $\text{rel}_{t_p^1}$
 513 transition in u_p^2 . Moreover if $ds(p)$ is green then we know that there is no $\text{rel}_{t_p^2}$ transition
 514 in u_p^2 .

515 We now combine these runs to get a run respecting σ ending in a deadlock configuration.
 516 For each process p such that $ds(p) = \perp$, execute the local run u_p . Since u_p is neutral, all
 517 locks are available after executing it. The only possible actions of p after this run are to
 518 acquire some locks from BT .

519 Next, for every process p such that $ds(p)$ is a green edge, execute the local run u_p^1 . This
 520 is also a neutral run. After this run p is in a state where σ_p allows to take lock t_p^1 , but p
 521 does not own any lock.

522 Next, in increasing order according to \leq , for every lock t with an outgoing blue edge
 523 $ds(p) = t \xrightarrow{p} t'$ execute the run u_p , except for the last $\text{acq}_{t'}$ action. After this run lock t is
 524 taken by p , and all actions allowed by σ_p are $\text{acq}_{t'}$ actions. Since there is only one outgoing
 525 edge from every lock, and since we are respecting the order \leq , both t and t' are free before
 526 executing that run. Hence it is possible to execute this run.

527 Finally, we come back to processes p such that $ds(p)$ is a green edge. For every such
 528 process we execute $\text{acq}_{t_p^1}$ followed by u_p^2 . This is possible because t_p^1 is free as there is a
 529 unique outgoing edge from t_p^1 . After executing these runs every process p with $ds(p) \neq \perp$ is
 530 in a state when the only possible action is $\text{acq}_{t_p^2}$.

531 At this stage all locks that are sources of edges from $ds(Proc)$ are taken. Since every lock
 532 in BT is a source of an edge, all locks from BT are taken. Thus no process p with $ds(p) = \perp$
 533 can move as it needs some lock from BT . Similarly, no process p with $ds(p) \neq \perp$ can move,
 534 as they need locks pointed by targets of the edges $ds(p)$, and these are in BT too. So we
 535 have constructed a run respecting σ and reaching a deadlock. ◀

536 From now on we concentrate on deciding if there is some deadlock scheme for a given
 537 graph G along with a lockset family $Locks$. Our approach will be to repeatedly eliminate
 538 edges from G or add locks to Z , and construct a deadlock scheme on Z at the same time.

539 As a preparatory step we observe that we can almost ignore the lockset family. Examining
 540 the definition of Z -deadlock scheme we see that the only information about $Locks$ it uses is
 541 whether $L_p = \emptyset$ or not. Hence we call a process *solid* if $L_p = \emptyset$, and *fragile* otherwise. The
 542 second condition in the definition of Z -deadlock scheme becomes: if $p \in Proc_Z$ is solid then
 543 $ds_Z(p) \neq \perp$.

544 The next lemma gives an important composition principle for deadlock schemes. Suppose
 545 we already have a set of “kernel” locks Z on which we know how to construct a Z -deadlock
 546 scheme. Then the lemma says that in order to get a deadlock scheme for G it is enough to
 547 consider the remaining part $G \setminus Z$.

548 ▶ **Lemma 23.** *Let $Z \subseteq T$ be such that there is no edge labeled by a solid process from a lock*
 549 *of Z to a lock of $T \setminus Z$ in G . Suppose $ds_Z : Proc_Z \rightarrow E \cup \{\perp\}$ is a Z -deadlock scheme. Then*
 550 *there is a deadlock scheme for G if and only if there is one equal to ds_Z over $Proc_Z$.*

551 The rest of the proof is a sequence of stages. We start with $H = G$ and $Z = \emptyset$. At each
 552 stage we remove some edges in H or extend Z . This process continues till some obstacle to
 553 the existence of a deadlock scheme is found, or till Z is big enough to be a deadlock scheme.
 554 We use three invariants:

555 ▶ **Invariant 1.** *G admits a deadlock scheme if and only if H does.*

556 ▶ **Invariant 2.** *There are no edges labeled by a solid process from Z to $T \setminus Z$ in H .*

557 ► **Invariant 3.** *There exists a Z -deadlock scheme.*

558 The relatively long proof of the following proposition is presented in Appendix B.2.

559 ► **Proposition 24.** *There is a polynomial time algorithm to decide if a lock graph G and a*
560 *lockset family $Locks$ have a deadlock scheme.*

561 The final argument behind Theorem 8 is as follows. We start by non-deterministically
562 guessing G and $Locks$. These are of polynomial size with respect to the size of the 2LSS. We
563 can check in polynomial time that there exists a strategy σ giving G and $Locks$ (Lemma 19).
564 If that is not the case, we reject the input. Otherwise we check if G and $Locks$ admit a
565 deadlock scheme (Proposition 24). By Lemma 22, the strategy σ is winning if and only if
566 the check says that there is no deadlock scheme in G and $Locks$.

567 **5 Solving the exclusive case in PTIME**

568 In this section we study exclusive 2LSS. We have shown an NP algorithm for the deadlock
569 avoidance control problem when restricting to locally live strategies. Here we show that the
570 problem is in PTIME if the 2LSS is exclusive (Definition 9). This is possible because the
571 exclusive assumption simplifies the structure of lock graphs, and makes the lockset family
572 unnecessary.

573 Throughout this section we fix an exclusive 2LSS, call it \mathcal{S} . The exclusive property
574 prohibits situations as in Figure 2 where a state has one outgoing $\text{acq}_{t_{p+1}}$ transition, and
575 one rel_{t_p} transition. Compared to the previous section we do not need to make a difference
576 between solid and fragile processes. We can even ignore colors on the arrows. This is a
577 consequence of the following two lemmas.

578 ► **Lemma 25.** *Let σ be a locally live control strategy and G_σ its lock graph. For all $t_1, t_2 \in T$,*
579 *if G_σ has a blue edge $t_1 \xrightarrow{p} t_2$ then it has a green edge $t_2 \xrightarrow{p} t_1$.*

580 ► **Lemma 26.** *Let σ be a locally live control strategy and G_σ its lock graph. For every edge*
581 *$t_1 \xrightarrow{p} t_2$ in G , process p is $\{t_1, t_2\}$ -lockable.*

582 Thanks to these simplifications there is a much more direct way of checking if a strategy
583 is winning. Take a locally live strategy σ . Consider a decomposition of G_σ into strongly
584 connected components (SCC). We say that an SCC is a *direct deadlock* if it contains at least
585 two nodes, and:

- 586 ■ either it has an edge that is not a double edge: $t_1 \xrightarrow{p} t_2$ but not $t_1 \xleftarrow{p} t_2$, for some p ;
- 587 ■ or all edges in the component are double edges and there is a proper cycle, i.e., all edges
588 are labeled by different processes.

589 A *deadlock SCC* is a direct deadlock SCC or an SCC that can reach some direct deadlock
590 SCC. Let BT_σ be the set of all the locks appearing in some deadlock SCC. We obtain a
591 simple characterization of winning strategies.

592 ► **Proposition 27.** *A strategy σ is winning if and only if there exists a process that is not*
593 *BT_σ -lockable.*

594 Building on this result we can give a method to decide if there is a winning strategy in
595 the system \mathcal{S} . For every process p and every set of edges between two locks of p we check
596 if there is a local strategy inducing exactly these edges. This can be done in a similar way
597 as Lemma 19. We say that an edge labeled by p is *unavoidable* if all the local strategies

598 σ_p induce this edge. Let G_S be the graph whose nodes are locks and edges are unavoidable
599 edges.

600 We calculate a set BT_S in a similar way as BT_σ in the previous proposition except that
601 we use slightly more general basic SCCs of G_S . A *direct semi-deadlock SCC* is either a direct
602 deadlock SCC or an SCC containing at least two nodes, only double edges, and two locks t_1
603 and t_2 such that for some process p not inducing a double edge between t_1, t_2 in G_S : every
604 strategy for p induces at least one edge between t_1 and t_2 . Then a *semi-deadlock SCC* is an
605 SCC that can reach some direct semi-deadlock SCC, or is itself a direct semi-deadlock SCC.

606 Let BT_S be the set of locks appearing in semi-deadlock SCCs of G_S . Theorem 11 follows
607 from the next proposition.

608 ► **Proposition 28.** *Let S be an exclusive 2LSS. There is a winning locally live strategy for*
609 *the system if and only if there exists a locally live strategy σ_p for some process p preventing*
610 *it from acquiring any lock from BT_S .*

611 The algorithm computes BT_S , and then checks if for some process p the condition from
612 the proposition holds. This check amounts to solving a safety game on a finite graph – the
613 transition graph of process p . The complete proof is presented in Appendix C.

614 **6 Nested-locking strategies**

615 We switch to another decidable case, where we require that locks are acquired and released in
616 stack-like manner. Our goal is Theorem 13 saying that the deadlock avoidance control problem
617 is NEXPTIME-complete when restricted to nested-locking strategies (cf. Definition 12).

618 In the context of this section we cannot assume that a process knows which locks it has
619 (cf. Remark 4). In consequence, it is not realistic to require that a strategy is locally live.
620 Yet, the lower bound works also for locally live strategies.

621 We will use some notions about local runs as defined on page 10.

622 ► **Definition 29.** *A stair decomposition of a local run u is*

$$623 \quad u = u_1 \mathbf{acq}_{t_1} u_2 \mathbf{acq}_{t_2} \dots u_k \mathbf{acq}_{t_k} u_{k+1}$$

624 *where in the configuration reached by $u_1 \mathbf{acq}_{t_1} u_2 \mathbf{acq}_{t_2} \dots u_i$ the set of locks held by the process*
625 *is $\{t_1, \dots, t_{i-1}\}$ for every $i > 0$, and there is no operation on t_i in $u_{i+1} \dots u_{k+1}$. (We omit*
626 *the actions associated with each operation as they are irrelevant here).*

627 Every nested-locking run has a unique stair decomposition.

628 Without the locally live assumption we may have runs simply ending because there are
629 no outgoing actions. Recall that given a strategy σ , a *risky σ -run* is a local σ -run ending in a
630 state from which every outgoing action allowed by σ acquires some lock. We define patterns
631 of risky local runs that will serve as witnesses of reachable deadlocks.

632 ► **Definition 30.** *Consider a stair decomposition $u_1 \mathbf{acq}_{t_1} u_2 \mathbf{acq}_{t_2} \dots u_k \mathbf{acq}_{t_k} u_{k+1}$ of a risky*
633 *σ -run u of a process p . Suppose the run is T_{blocks} -blocked, and let $T_{\text{owns}} = \{t_1, \dots, t_k\}$. We*
634 *associate with u a stair pattern $(T_{\text{owns}}, T_{\text{blocks}}, \preceq)$, where \preceq is the smallest partial order on*
635 *the set T_p of locks of p satisfying: for all i , for all $t \in T_p$, if the last operation on t in the*
636 *run is after the last \mathbf{acq}_{t_i} then $t_i \preceq t$. A behavior of σ is a family of sets of stair patterns*
637 *$(\mathcal{P}_p)_{p \in \text{Proc}}$, where \mathcal{P}_p is the set of stair patterns of local risky σ -runs of p .*

638 Similarly to Lemma 22 we can show that the family of patterns for a strategy determines
639 if it is winning.

640 ► **Lemma 31.** *A nested-locking control strategy σ with behavior $(\mathcal{P}_p)_{p \in Proc}$ is **not** winning*
 641 *if and only if for every $p \in Proc$ there is a stair pattern $(T_{owns}^p, T_{blocks}^p, \preceq^p) \in \mathcal{P}_p$ such that:*

642 ■ $\bigcup_{p \in Proc} T_{blocks}^p \subseteq \bigcup_{p \in Proc} T_{owns}^p,$

643 ■ *the sets T_{owns}^p are pairwise disjoint,*

644 ■ *there exists a total order \preceq , on the set of all locks T , compatible with all \preceq^p .*

645 Similarly to Lemma 19 we can check if there is a strategy whose set of patterns has only
 646 patterns from a given family. Observe that the depth of nesting is bounded by the number
 647 of locks.

648 ► **Lemma 32.** *Given a lock-sharing system $((\mathcal{A}_p)_{p \in Proc}, \Sigma^s, \Sigma^e, T)$, a process $p \in Proc$ and*
 649 *a set of patterns \mathcal{P}_p , we can check in polynomial time in $|\mathcal{A}_p|$ and $2^{|T|}$ whether there exists a*
 650 *nested-locking local strategy σ_p with set of patterns included in \mathcal{P}_p .*

651 ► **Proposition 33.** *The deadlock avoidance control problem is decidable for lock-sharing*
 652 *systems with nested-locking strategies in non-deterministic exponential time.*

653 **Proof.** The decision procedure guesses a set of patterns \mathcal{P}_p for each process p , of size at
 654 most $2^{2^{|T|}}|T|! \leq 2^{O(|T| \log(|T|))}$. Then it checks if there exist local strategies yielding subsets
 655 of those sets of patterns. This takes exponential time by Lemma 32. If the result is negative
 656 then the procedure rejects. Otherwise, it checks if some condition from Lemma 31 does not
 657 hold. If it finds one then it accepts, otherwise it rejects.

658 Clearly, if there is a winning nested-locking strategy then the procedure can accept by
 659 guessing the family of patterns corresponding to this strategy. For this family the check from
 660 Lemma 32 does not fail, and one of the conditions of Lemma 31 must be violated.

661 Conversely, if the decision procedure concludes that there exists a winning strategy, then
 662 let $(\mathcal{P}_p)_{p \in Proc}$ be the guessed family of sets of patterns. We know that there exists a strategy
 663 σ with behaviors $(\mathcal{P}'_p)_{p \in Proc}$ such that $\mathcal{P}'_p \subseteq \mathcal{P}_p$ for all $p \in Proc$. Furthermore, as there are
 664 no patterns in $(\mathcal{P}_p)_{p \in Proc}$ satisfying the requirements of Lemma 31, there cannot be any in
 665 the \mathcal{P}'_p either. Hence σ is a winning strategy. ◀

666 The proof of the matching lower bound and the missing lemmas are in Appendix D.

667 **7 Undecidability for unrestricted lock-sharing systems**

668 In this section we show that the deadlock avoidance control problem for lock-sharing systems
 669 is undecidable for three processes with a fixed number of locks. Three locks used in non-nested
 670 fashion allow to synchronize two processes in lock-step manner. This is an essential ingredient
 671 for the undecidability proof.

672 We have defined lock-sharing systems so that initially all locks are free. First we show the
 673 undecidability result supposing that we are allowed to start with a designated distribution of
 674 locks. Later we describe how to implement initial lock distributions using extra locks.

675 ► **Lemma 34.** *The control problem for lock-sharing systems with 3 processes, fixed initial*
 676 *configuration and fixed number of locks per process is undecidable.*

677 The proof uses the usual recipe for the undecidability of distributed synthesis [26, 27].
 678 Two processes P and \bar{P} synchronize with a third process C over a stream of bits chosen
 679 by their strategy. The process C is partially controlled by the environment, which selects
 680 non-deterministically an interleaving of the two streams and parses the interleaving with a
 681 finite automaton. This is enough to get undecidability by a reduction from an infinite Post
 682 Correspondence Problem (PCP).

683 Consider an instance $(\alpha_i, \beta_i)_{i \in I}$ of PCP on the alphabet $\{0, 1\}$. A solution is an infinite
 684 sequence $i_1 i_2 \dots \in I^\omega$ such that $\alpha_{i_1} \alpha_{i_2} \dots = \beta_{i_1} \beta_{i_2} \dots$. The two streams sent by P and \bar{P}
 685 to C , are $\alpha = \alpha_{i_1} i_1 \alpha_{i_2} i_2 \dots$ and $\beta = \beta_{j_1} j_1 \beta_{j_2} j_2 \dots$, resp. With finite memory C can check
 686 equality of the two words ($\alpha_{i_1} \alpha_{i_2} \dots = \beta_{j_1} \beta_{j_2} \dots$) or equality of the two index sequences
 687 ($i_1 i_2 \dots = j_1 j_2 \dots$). Since P and \bar{P} are not aware of what C does, the streams are fixed by
 688 the strategies and do not depend on what C is checking.

689 The locks used in the proof are $\{c, s_0, s_1, p, \bar{c}, \bar{s}_0, \bar{s}_1, \bar{p}\}$. Process C and P use locks from
 690 $\{c, s_0, s_1, p\}$ to synchronize and similarly for C, \bar{P} and $\{\bar{c}, \bar{s}_0, \bar{s}_1, \bar{p}\}$.

691 It remains to explain the synchronization mechanism. The two processes P and C
 692 synchronize over a bit of information, say bit 0, by executing specific finite runs using the
 693 locks $\{s_0, c, p\}$ in non-nested fashion. Initially, C owns $\{s_0, c\}$ and P owns $\{p\}$. First, C
 694 releases lock s_0 and P acquires it, which we denote as $C \xrightarrow{s_0} P$. Here, P is waiting for C
 695 to release s_0 , and the two actions rel_{s_0} of C and acq_{s_0} of P are ordered. The rest of the
 696 run follows a similar pattern: at each step, one of the processes is waiting to take a lock
 697 released by the other process. With the same notation, the run proceeds with $P \xrightarrow{p} C$, and
 698 continues until each process owns the same locks it owned at the start: each lock is sent
 699 twice, from its initial owner to the other process, and back. To sum up, the exchange of bit
 700 0 between C and P corresponds to:

$$701 \quad C \xrightarrow{s_0} P \xrightarrow{p} C \xrightarrow{c} P \xrightarrow{s_0} C \xrightarrow{p} P \xrightarrow{c} C .$$

702
 703 In other words, processes C and P respectively perform two local runs:

$$704 \quad C : \text{rel}_{s_0} \text{acq}_p \text{rel}_c \text{acq}_{s_0} \text{rel}_p \text{acq}_c \quad P : \text{acq}_{s_0} \text{rel}_p \text{acq}_c \text{rel}_{s_0} \text{acq}_p \text{rel}_c$$

705 Observe that P and C need to execute these sequences in lock-step manner, as one of the
 706 two processes waits for a lock from the other.

707 In order to synchronize over bit 1, the two processes perform a similar synchronization,
 708 using s_1 instead of s_0 . The communication between C and \bar{P} is identical, except that it uses
 709 locks from $\{\bar{c}, \bar{s}_0, \bar{s}_1, \bar{p}\}$.

710 In each round, P and C must agree beforehand on a bit they are going to synchronize on,
 711 either s_0 or s_1 . Otherwise the two processes get blocked, and \bar{P} will get blocked too, as it
 712 needs locks held by C .

713 A bit stream between C and P is encoded as a concatenation of such runs, and similarly
 714 for C, \bar{P} . The content of the two bit streams is chosen by the strategies of P, \bar{P}, C . Since the
 715 strategy has infinite memory, there is no upper bound on the complexity of the streams.

716 Interestingly, two locks are not enough for two processes to synchronize over a bit stream.

717 The next lemma shows a generic reduction of the control problem with initial configuration
 718 to the one where all locks are initially free.

719 **► Lemma 35.** *There is a polynomial-time reduction from the control problem for lock-sharing*
 720 *systems with initial configuration to the control problem where all locks are initially free. The*
 721 *reduction adds $|Proc|$ new locks.*

722 We sketch the proof idea. Assume that we have pairwise disjoint sets $(I_p)_{p \in Proc}$ of locks,
 723 and a lock-sharing system \mathcal{S} in which each process p initially owns exactly the locks in I_p .
 724 We build another lock-sharing systems \mathcal{S}_\emptyset that starts with all locks initially free, makes every
 725 process acquire all locks in I_p , and then simulates \mathcal{S} .

726 It is important that the initialization phase of \mathcal{S}_\emptyset does not interfere with the simulation
 727 of \mathcal{S} . We ensure this by using one additional lock k_p per process, called the “key” of p .

728 For process p , the initialization sequence consists of three steps.

- 729 1. First, p takes one by one (in a fixed arbitrary order), all its initial locks in I_p .
 730 2. Second, p takes and releases, one by one (in a fixed arbitrary order) all the keys of the
 731 other processes $(k_q)_{q \neq p}$.
 732 3. Finally, p acquires its key k_p and keeps it forever.
 733 After acquiring k_p process p reaches the initial state in \mathcal{S} .

734 In order to prevent the initialization phase to create extra deadlocks, there is a local *nop*
 735 loop on every state of the initialization sequence. This way, a deadlock may only occur if all
 736 processes have finally completed their initialization sequences.

737 Let us explain why the initialization phase does not interfere with the simulation of \mathcal{S} .
 738 The exchange of keys guarantees that up to the moment where a process p has completed
 739 the initialization in \mathcal{S}_\emptyset , no other process has used any lock from I_p . The details of the
 740 construction are presented in Appendix E.

741 8 Conclusions

742 Motivated by a recent undecidability result for distributed control synthesis [17] we have
 743 considered a model for which the problem has not been investigated yet. With hindsight it
 744 is strange that the well-studied model of lock synchronization has not been considered in
 745 the context of distributed synthesis. One reason may be the "non-monotone" nature of the
 746 synthesis problem. It is not the case that for a less expressive class of systems the problem is
 747 necessarily easier because the controllers get less powerful, too.

748 The two decidable classes of lock-sharing systems presented here are rather promising.
 749 Especially because the low complexity results cover already non-trivial problems. All our
 750 algorithms are based on analyzing lock patterns. While in this paper we consider only finite
 751 state processes, the same method applies to more complex systems, as long as solving the
 752 centralized control problem in the style of Lemma 19 is decidable. This is for example the
 753 case for pushdown systems.

754 There are numerous directions that need to be investigated further. We have focused on
 755 deadlock avoidance because this is a central property, and deadlocks are difficult to discover
 756 by means of testing or verification. Another option is partial deadlock, where some, but not
 757 all, processes are blocked. The concept of Z -deadlock scheme from Definition 20 should help
 758 here, but the complexity results may be different. Reachability, and repeated reachability
 759 properties need to be investigated, too.

760 We do not know if the upper bound from Theorem 8 is tight. The algorithm for verifying
 761 if there is a deadlock in a given strategy graph, Proposition 24, is already quite complicated,
 762 and it is not clear how to proceed when a strategy is not given.

763 Another research direction is to consider probabilistic controllers. It is well known that
 764 there are no symmetric solutions to the dining philosophers problem but there is a randomized
 765 one [21, 22]. Symmetric solutions are quite important for resilience issues as it is preferable
 766 that every process runs the same code. The Lehmann-Rabin algorithm is essentially the
 767 system presented in Figure 2 where the choice between *left* and *right* is made randomly. This
 768 is one of the examples where randomized strategies are essential. Distributed synthesis has a
 769 potential here because it is even more difficult to construct distributed randomized systems
 770 and prove them correct.

References

- 771 ——— **References** ———
- 772 1 A. Arnold and I. Walukiewicz. Nondeterministic controllers of nondeterministic processes. In
773 Jörg Flum, Erich Grädel, and Thomas Wilke, editors, *Logic and Automata*, volume 2 of *Texts*
774 *in Logic and Games*, pages 29–52. Amsterdam University Press, 2007.
- 775 2 B. Bérard, B. Bollig, P. Bouyer, M. Függer, and N. Sznajder. Synthesis in presence of
776 dynamic links. In Jean-François Raskin and Davide Bresolin, editors, *Proceedings 11th*
777 *International Symposium on Games, Automata, Logics, and Formal Verification, GandALF*
778 *2020*, volume 326 of *EPTCS*, pages 33–49, 2020. To appear in *Information and Computation*.
779 doi:10.4204/EPTCS.326.3.
- 780 3 R. Beutner, B. Finkbeiner, and J. Hecking-Harbusch. Translating asynchronous games for
781 distributed synthesis. In *International Conference on Concurrency Theory (CONCUR'19)*,
782 volume 140 of *LIPIcs*, pages 26:1–26:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik,
783 2019.
- 784 4 K. M. Chandy and J. Misra. The drinking philosophers problem. *ACM Trans. Program. Lang.*
785 *Syst.*, 6(4):632–646, oct 1984. doi:10.1145/1780.1804.
- 786 5 A. Church. Applications of recursive arithmetic to the problem of circuit synthesis. In
787 *Summaries of the Summer Institute of Symbolic Logic*, volume I, pages 3–50. Cornell Univ.,
788 Ithaca, N.Y., 1957.
- 789 6 E. M. Clarke and E. A. Emerson. Design and synthesis of synchronization skeletons using
790 branching time temporal logic. In *Workshop on Logics of Programs*, volume 131 of *Lecture*
791 *Notes in Computer Science*, pages 52–71. Springer Verlag, 1981.
- 792 7 M. D. Ernst, A. Lovato, D. Macedonio, F. Spoto, and J. Thaine. Locking discipline inference
793 and checking. In *ICSE 2016, Proceedings of the 38th International Conference on Software*
794 *Engineering*, pages 1133–1144, Austin, TX, USA, May 2016.
- 795 8 B. Finkbeiner. Bounded synthesis for Petri games. In Roland Meyer, André Platzer, and
796 Heike Wehrheim, editors, *Correct System Design - Symposium in Honor of Ernst-Rüdiger*
797 *Olderog on the Occasion of His 60th Birthday, Oldenburg, Germany, September 8-9, 2015.*
798 *Proceedings*, volume 9360 of *Lecture Notes in Computer Science*, pages 223–237. Springer,
799 2015. doi:10.1007/978-3-319-23506-6_15.
- 800 9 B. Finkbeiner, M. Giesekeing, J. Hecking-Harbusch, and E.-R. Olderog. Global winning
801 conditions in synthesis of distributed systems with causal memory. In Florin Manea and Alex
802 Simpson, editors, *30th EACSL Annual Conference on Computer Science Logic, CSL 2022,*
803 *Virtual Conference*, volume 216 of *LIPIcs*, pages 20:1–20:19. Schloss Dagstuhl - Leibniz-Zentrum
804 für Informatik, 2022. doi:10.4230/LIPIcs.CSL.2022.20.
- 805 10 B. Finkbeiner and E.-R. Olderog. Petri games: Synthesis of distributed systems with causal
806 memory. *Inf. Comput.*, 253:181–203, 2017.
- 807 11 B. Finkbeiner and S. Schewe. Uniform distributed synthesis. In *LICS'05*, pages 321–330. IEEE
808 Computer Society, 2005.
- 809 12 P. Gastin, B. Lerman, and M. Zeitoun. Distributed games with causal memory are decidable
810 for series-parallel systems. In *FSTTCS'04*, volume 3328 of *LNCS*, pages 275–286. Springer,
811 2004.
- 812 13 P. Gastin, N. Sznajder, and M. Zeitoun. Distributed synthesis for well-connected architectures.
813 *Formal Methods in System Design*, 34(3):215–237, June 2009.
- 814 14 B. Genest, H. Gimbert, A. Muscholl, and I. Walukiewicz. Asynchronous games over tree archi-
815 tectures. In *International Colloquium on Automata, Languages and Programming (ICALP'13)*,
816 volume 7966 of *LNCS*, pages 275–286. Springer, 2013.
- 817 15 M. Giesekeing, J. Hecking-Harbusch, and A. Yanich. A web interface for Petri nets with
818 transits and Petri games. In Jan Friso Groote and Kim Guldstrand Larsen, editors, *Tools*
819 *and Algorithms for the Construction and Analysis of Systems - 27th International Conference,*
820 *TACAS 2021, Held as Part of the European Joint Conferences on Theory and Practice of*
821 *Software, ETAPS 2021, Proceedings, Part II*, volume 12652 of *Lecture Notes in Computer*
822 *Science*, pages 381–388. Springer, 2021. doi:10.1007/978-3-030-72013-1_22.

- 823 16 H. Gimbert. On the control of asynchronous automata. In *FSTTCS'17*, volume 30 of *LIPICs*.
824 Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- 825 17 H. Gimbert. Distributed asynchronous games with causal memory are undecidable. *CoRR*,
826 abs/2110.14768, 2021. Submitted. URL: <https://arxiv.org/abs/2110.14768>, arXiv:2110.
827 14768.
- 828 18 J. Hecking-Harbusch and N. O. Metzger. Efficient trace encodings of bounded synthesis for
829 asynchronous distributed systems. In Yu-Fang Chen, Chih-Hong Cheng, and Javier Esparza,
830 editors, *Automated Technology for Verification and Analysis - 17th International Symposium*,
831 *ATVA 2019, Proceedings*, volume 11781 of *Lecture Notes in Computer Science*, pages 369–386.
832 Springer, 2019. doi:10.1007/978-3-030-31784-3_22.
- 833 19 V. Kahlon and A. Gupta. An automata-theoretic approach for model checking threads for
834 LTL properties. In *21st Annual IEEE Symposium on Logic in Computer Science (LICS'06)*,
835 pages 101–110, 2006. doi:10.1109/LICS.2006.11.
- 836 20 O. Kupferman and M. Y. Vardi. Synthesizing distributed systems. In *LICS'01*, pages 389–398.
837 IEEE, 2001.
- 838 21 D. Lehmann and M. O. Rabin. On the advantages of free choice: A symmetric and fully
839 distributed solution to the dining philosophers problem. In John White, Richard J. Lipton,
840 and Patricia C. Goldberg, editors, *Conference Record of the Eighth Annual ACM Symposium*
841 *on Principles of Programming Languages, Williamsburg, Virginia, USA, January 1981*, pages
842 133–138. ACM Press, 1981. doi:10.1145/567532.567547.
- 843 22 N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.
- 844 23 P. Madhusudan, P. S. Thiagarajan, and S. Yang. The MSO theory of connectedly communic-
845 ating processes. In *FSTTCS'05*, volume 3821 of *LNCS*, pages 201–212. Springer, 2005.
- 846 24 P. Madhusudan and P.S. Thiagarajan. Distributed control and synthesis for local specifications.
847 In *ICALP'01*, volume 2076 of *LNCS*, pages 396–407. Springer, 2001.
- 848 25 A. Muscholl and I. Walukiewicz. Distributed synthesis for acyclic architectures. In *FSTTCS'14*,
849 volume 29 of *LIPICs*, pages 639–651. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik,
850 2014.
- 851 26 G. L. Peterson and J. H. Reif. Multiple-person alternation. In *20th Annual Symposium on*
852 *Foundations of Computer Science (SFCS 1979)*, pages 348–363. IEEE, 1979.
- 853 27 A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. ACM POPL*, pages
854 179–190, 1989.
- 855 28 A. Pnueli and R. Rosner. Distributed reactive systems are hard to synthesize. In *FOCS'90*,
856 pages 746–757. IEEE Computer Society, 1990.
- 857 29 P. J.G. Ramadge and W. M. Wonham. The control of discrete event systems. *Proceedings of*
858 *the IEEE*, 77(2):81–98, 1989.
- 859 30 K. Rudie and W. M. Wonham. Think globally, act locally: Decentralized supervisory control.
860 *IEEE Trans. on Automat. Control*, 37(11):1692–1708, 1992.
- 861 31 J. G. Thistle. Undecidability in decentralized supervision. *Systems & Control Letters*, 54(5):503–
862 509, 2005.
- 863 32 S. Tripakis. Undecidable problems in decentralized observation and control for regular
864 languages. *Information Processing Letters*, 90(1):21–28, 2004.
- 865 33 I. Walukiewicz. Synthesis with finite automata. In J. E. Pin, editor, *Handbook of Automata*
866 *Theory*, volume 2, pages 1215–1258. 2021. [https://www.labri.fr/perso/igw/Papers/igw-synt-](https://www.labri.fr/perso/igw/Papers/igw-synt-chapter.pdf)
867 [chapter.pdf](https://www.labri.fr/perso/igw/Papers/igw-synt-chapter.pdf).
- 868 34 Y. Wang, S. Lafortune, T. Kelly, M. Kudlur, and S. A. Mahlke. The theory of deadlock
869 avoidance via discrete control. In Zhong Shao and Benjamin C. Pierce, editors, *Proceedings*
870 *of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*,
871 *POPL 2009*, pages 252–263. ACM, 2009. doi:10.1145/1480881.1480913.

872 **A** Two locks per process: Σ_2^P -completeness

873 In this section we prove Theorem 7, as recalled below:

874 ► **Theorem 7.** *The deadlock avoidance control problem for 2LSS is Σ_2^P -complete.*

875 Thanks to Lemma 15, in order to decide if there is a winning strategy for a given system
876 we need to come up with a set of patterns $Patt_p$ for each process p and show two things:

- 877 ■ these sets of patterns do not meet the conditions given in Lemma 15.
- 878 ■ there is a strategy σ whose local runs on each process p all match a pattern of $Patt_p$.

879 Note that in the second condition we only require an inclusion because it is clear from
880 the previous lemma that the less patterns a strategy allows, the less chances there are it
881 leads to a deadlock.

882 We start by showing that we can check the second condition in polynomial time.

883 ► **Lemma 36.** *Given a family $(Patt_p)_{p \in Proc}$ of sets of patterns, it is decidable in PTIME
884 whether there exists a strategy $\sigma = (\sigma_p)_{p \in Proc}$ such that for all p and all σ_p -runs u_p of p , the
885 pattern of u_p is in $Patt_p$.*

886 **Proof.** First of all note that we only need to check for each p that there exists a local strategy
887 σ_p that does not allow any risky runs whose pattern is not in $Patt_p$.

888 Let $p \in Proc$, let \mathcal{A}_p be its transition system. We extend it in a similar way as in
889 Remark 4, by adding some information in the states. We already assumed that the states
890 contained the information of which locks are currently owned by p . We duplicate the states
891 where p only owns one lock t_1 , in order to add a bit of information saying whether p released
892 its other lock t_2 since acquiring t_1 for the last time.

893 This way, the risky nature of local runs and their patterns depend only on the state
894 in which they end and its outgoing transitions. For instance if a state has no outgoing
895 transitions and is such that when reaching it p holds t_1 and released t_2 since acquiring it,
896 then the pattern of runs ending there is $(\{t_1\}, \emptyset, (t_1, t_2))$. If this pattern is not in $Patt_p$ then
897 we declare this state as bad.

898 Formally, we define as good all states such that there exists a set of outgoing transitions
899 containing all environment transitions and such that

- 900 ■ either it contains a transition with no acquire operation
- 901 ■ or the set B of locks gotten by those transitions is such that runs ending in that state
902 have a pattern in $Patt_p$.

903 We define the other states as bad. If all states of the system have that property then
904 clearly there is a suitable strategy.

905 Otherwise a local control strategy cannot allow any run to reach a bad state without
906 getting a pattern outside of the input set, hence the following algorithm. It resembles the
907 usual algorithm for solving safety games, except that we do not simply want to avoid some
908 states, but we want to avoid having to allow some sets of actions from some states.

909 We simply iteratively delete bad states and all their ingoing transitions. If one of those
910 transitions is controlled by the environment, we declare its source state as bad (as reaching
911 that state would allow the environment to take that transition, leading us to a bad state).
912 Note that deleting transitions may create more bad states by reducing the choice of the
913 system. If we end up deleting all states, we conclude that there is no suitable strategy.
914 Otherwise the subsystem we obtain only has good states, allowing us to get a strategy
915 matching the input set of patterns. ◀

916 ► **Proposition 37.** *The deadlock avoidance control problem for 2LSS is decidable in Σ_2^P .*

917 **Proof.** We first non-deterministically guess a set of patterns $Patt_p$ for each process p (each
918 set is of bounded size). By Lemma 36, we can then check in polynomial time if there exists a
919 strategy σ respecting that set of patterns.

920 If it exists, then we have an adversarial selection of a pattern $pat_p \in Patt_p$ for each p , as
921 well as an adversarial guess of a total order on locks \leq . It is then easy to check in polynomial
922 time if these patterns meet the conditions of Lemma 15.

923 If they do, we reject the input, otherwise we accept it.

924 If there exists a winning strategy then we take the sets of patterns it allows, we conclude
925 that the system wins by Lemma 15.

926 Conversely if we find sets of patterns not meeting the requirements of Lemma 15 and
927 such that there is a strategy σ respecting them, then the sets of patterns allowed by σ are
928 subsets of the ones we guessed and therefore they do not meet the conditions of Lemma 15
929 either. Hence σ is a winning strategy. ◀

930 We provide the matching lower bound.

931 ► **Proposition 38.** *The deadlock avoidance control problem for 2LSS is Σ_2^P -hard, even when
932 restricted to exclusive systems.*

933 **Proof.** We reduce from the $\exists\forall$ -SAT problem. We are given a boolean formula in 3-
934 disjunctive normal form $\bigvee_{i=1}^k \alpha_i$; each clause α_i is a conjunction of three literals $\ell_1^i \wedge \ell_2^i \wedge \ell_3^i$
935 over a set of variables $\{x_1, \dots, x_n, y_1, \dots, y_m\}$. The question is whether the formula
936 $\exists x_1, \dots, x_n \forall y_1, \dots, y_m, \bigvee_{i=1}^k \alpha_i$ is true.

937 We construct a 2LSS for which there is a winning strategy iff the formula is true. The
938 2LSS will use locks:

$$939 \quad \{t_i \mid 1 \leq i \leq k\} \cup \{x_i, \bar{x}_i \mid 1 \leq i \leq n\} \cup \{y_j, \bar{y}_j \mid 1 \leq j \leq m\} .$$


940 For all $1 \leq i \leq n$ we have a process p_i with four states, as depicted in Figure 3. In that
941 process the system has to take both x_i and \bar{x}_i , and then may release one of them before
942 being blocked in a state with no outgoing transitions. Similarly, for all $1 \leq j \leq m$ we have a
943 process q_j , in which the environment has to take y_j or \bar{y}_j , and then is blocked.

944 For each clause α_i we also have a process $p(\alpha_i)$ which just has one transition acquiring lock
945 t_i towards a state with a local loop on it. Hence to block all those processes the environment
946 needs to have all t_i taken by other processes.

947 It can do that with our last kind of processes. For each clause α_i and each literal ℓ of α_i
948 there is a process $p_i(\ell)$. There the environment has to acquire t_i and then ℓ before entering
949 a state with a self-loop.

950 The only way to block $p_i(\ell)$ is to have either the t_i or the ℓ lock taken by another process.
951 Intuitively, in the first case the environment accepts that the literal ℓ is true while in the
952 second case the environment claims that the literal ℓ is false and has to prove his claim.

953 A strategy for the system amounts to choosing whether p_i should release x_i or \bar{x}_i , for
954 each $i = 1, \dots, n$. It may also choose to release neither. Since the environment has a global
955 view of the system, it can afterwards choose one of y_j, \bar{y}_j in process q_j , for each $j = 1, \dots, m$.
956 Those choices represent valuations, the free lock remaining being the satisfied literal. In
957 order to win, the environment has to ensure that no lock t_j is free (otherwise some $p(C_j)$
958 would not be blocked), by choosing a literal ℓ_j whose corresponding lock is not free for every
959 formula α_j and taking t_j in $p_j(\ell_j)$. This means that the environment wins if and only if for



■ **Figure 3** The processes used in the reduction in Proposition 38. Transitions of the system are dashed.

960 all formulas α_i , one of the literals of α_i is taken. It is therefore never in the interest of the
 961 system to release neither x_i nor \bar{x}_i . We can assume that it always releases one of them.

962 Equivalently, the system wins if and only if there exists a valuation of the x 's such that
 963 for every valuation of the y 's there is at least one α_j whose three literals are all satisfied.
 964 This concludes our reduction.

965

966 **B Two locks per process with locally live strategies**

967 We provide missing proofs and constructions from Section 4.

968 **B.1 Characterization of winning strategies (Lemma 22)**

969 Lemma 22 states that the information given by the lock graph and lockset family of a locally
 970 live strategy is enough to determine if it is winning. Furthermore, it gives us a goal for the
 971 verification of such witnesses: We have to check whether there exists a deadlock scheme.

972 We prove the lemmas necessary to our polynomial-time algorithm checking the existence
 973 of a deadlock scheme.

974 ► **Lemma 23.** *Let $Z \subseteq T$ be such that there is no edge labeled by a solid process from a lock*
 975 *of Z to a lock of $T \setminus Z$ in G . Suppose $ds_Z : Proc_Z \rightarrow E \cup \{\perp\}$ is a Z -deadlock scheme. Then*
 976 *there is a deadlock scheme for G if and only if there is one equal to ds_Z over $Proc_Z$.*

977 **Proof.** Suppose (BT, ds) is a deadlock scheme for G_σ . We construct another one ds' which
 978 is equal to ds_Z over $Proc_Z$. For every process $p \in Proc$, we define $ds'(p)$ as:

- 979 ■ $ds_Z(p)$ if $p \in Proc_Z$,
- 980 ■ \perp if p labels an edge from Z to $T \setminus Z$,
- 981 ■ $ds(p)$ otherwise.

982 We assumed that edges from Z to $T \setminus Z$ could not be labeled by solid processes, thus all
 983 processes mapped to \perp are fragile. Every lock $t \in BT$ has at most one outgoing edge in ds' ,
 984 since it can only come from ds_Z , if $t \in Z$, or from ds , if $t \in BT \setminus Z$. We verify that there is
 985 at least one outgoing edge. By definition of Z -deadlock scheme there is one outgoing edge
 986 from every lock in Z . A lock $t \in BT \setminus Z$ has exactly one outgoing edge in $ds(Proc)$, and this
 987 edge stays in $ds'(p)$.

988 Finally, there cannot be any blue cycle in $ds'(Proc)$ as there are none within Z or $BT \setminus Z$
 989 and all edges between the two sets are towards Z . ◀

990 B.2 PTIME procedure to check the existence of a deadlock scheme

991 We recall the proposition we want to prove

992 ▶ **Proposition 24.** *There is a polynomial time algorithm to decide if a lock graph G and a*
 993 *lockset family $Locks$ have a deadlock scheme.*

994 We will describe several polynomial-time algorithms operating on graph $H = (T, GE)$ and
 995 a set Z of locks. Observe that H will always have all locks as nodes. Each of those algorithms
 996 will either eliminate some edges from H or extend Z , while maintaining Invariants 1–3,
 997 recalled below.

998 ▶ **Invariant 1.** *G admits a deadlock scheme if and only if H does.*

999 ▶ **Invariant 2.** *There are no edges labeled by a solid process from Z to $T \setminus Z$ in H .*

1000 ▶ **Invariant 3.** *There exists a Z -deadlock scheme.*

1001 We start with H being the given G_σ and $Z = \emptyset$. The invariants are clearly satisfied.

1002 There is a simplification we can make: for most of the algorithm we will not use all of
 1003 $Locks_\sigma$ but simply distinguish between *solid* processes that are not B -lockable for any B
 1004 (i.e., that will necessarily be mapped to an edge in a deadlock scheme) and the others (called
 1005 *fragile*).

1006 ▶ **Definition 39.** *A process p is called solid if $L_p = \emptyset$ and fragile otherwise.*

1007 ▶ **Definition 40** (Double and solo solid edges). *Consider a solid process p . We say that there*
 1008 *is a double solid edge $t_1 \xleftrightarrow{p} t_2$ in H if both $t_1 \xrightarrow{p} t_2$ and $t_1 \xleftarrow{p} t_2$ exist in H . We say that*
 1009 *$t_1 \xrightarrow{p} t_2$ in H is a solo solid edge if there is no $t_1 \xleftarrow{p} t_2$ in H .*

1010 Our first algorithm looks for a solo solid edge $t_1 \xrightarrow{p} t_2$ and erases all other outgoing edges
 1011 from t_1 . It is correct as a deadlock scheme for H has to map p to the edge $t_1 \xrightarrow{p} t_2$ and there
 1012 must be at most one outgoing edge from every lock.

■ **Algorithm 1** Trimming the graph:

```

1: Find  $t \in H \setminus Z$  with a solo solid edge  $t \xrightarrow{p} t' \in EH$  and some other outgoing edges
2: If there is no such edge then stop and report success.
3: for every edge  $t \xrightarrow{q} t' \in HE$  from  $t$  with  $q \neq p$  do
4:   if  $q$  is solid and  $t \xleftarrow{q} t' \notin HE$  then
5:     return "Fail: no  $H$ -deadlock scheme"
6:   else
7:     Erase  $t \xrightarrow{q} t'$ 
8:   end if
9: end for

```

1013 We repeat this algorithm till no edges are removed. If some call to of the algorithm fails
 1014 then there is no full deadlock scheme for H . Otherwise the resulting H satisfies the property:

1015 (Trim) if a lock t in $H \setminus Z$ has an outgoing solo solid edge then it has no other
 1016 outgoing edges.

1017 We call H *trimmed* if it satisfies property (Trim).

1018 ► **Lemma 41.** *Suppose (H, Z) satisfies Invariants 1–3. If Algorithm 1 fails then there is no*
 1019 *H -deadlock scheme. After a successful execution of the algorithm all the invariants are still*
 1020 *satisfied. If a successful execution does not remove an edge from H then H satisfies (Trim).*

1021 **Proof.** Let H' be the graph after an execution of Algorithm 1. Observe that the algorithm
 1022 does not change Z . If $H = H'$ then (Trim) holds. If the algorithm fails then there is a lock
 1023 with two solo solid outgoing edges. Thus it is impossible to find a H -deadlock scheme.

1024 Finally, if the algorithm succeeds but H' is smaller than H , we must show that all the
 1025 invariants hold. Since the algorithm does not change Z , Invariants 2 and 3 continue to hold.
 1026 For Invariant 1, suppose $t \xrightarrow{p} t'$ is the edge found by the algorithm. Observe that if ds_H is
 1027 an H -deadlock scheme then $ds_H(p)$ must be this edge. So ds_H is also a deadlock scheme for
 1028 H' . In the other direction, an H' -deadlock scheme is also an H -deadlock scheme as H' is
 1029 a subgraph of H and $Proc_H = Proc_{H'}$. The latter holds because H' has the same locks as
 1030 H . ◀

1031 Our second algorithm searches for cycles formed by solid edges and eventually adds them
 1032 to Z . If the found cycle has a green edge then it can be added to Z . If the cycle is blue, it
 1033 may still be the case that its reversal is green. More precisely it may be the case that for
 1034 every solid edge $t_i \xrightarrow{p_i} t_{i+1}$ in the cycle there is also a reverse edge $t_i \xleftarrow{p_i} t_{i+1}$ (and it is
 1035 solid by definition). If the reversed cycle is also blue then there is no H -deadlock scheme.
 1036 If it does, Z can be added to H . The result still satisfies the invariants thanks to (Trim)
 1037 property of H .

■ **Algorithm 2** Find solid cycles and add them to Z if possible.

```

1: Find a simple cycle of solid edges  $t_1 \xrightarrow{p_1} t_2 \cdots \xrightarrow{p_k} t_{k+1} = t_1$  not intersecting  $Z$ , all  $t_i$ 
   distinct
2: if there is no such cycle, stop and report success.
3: if all the edges on the cycle are blue then
4:   if for some  $j$  there is no reverse edge  $t_j \xleftarrow{p_j} t_{j+1} \in EH$  then
5:     return "Fail: no  $H$ -deadlock scheme"
6:   else if all edges  $t_j \xleftarrow{p_j} t_{j+1}$  are blue then
7:     return "Fail: no  $H$ -deadlock scheme"
8:   end if
9: end if
10:  $Z \leftarrow Z \cup \{t_1, \dots, t_k\}$ 
11: For every  $t_i$  remove from  $H$  all edges outgoing from  $t_i$  but those that are on the cycle.
12: if some solid process  $p$  has no edge in  $H$  then
13:   return "Fail: no  $H$ -deadlock scheme"
14: end if
15: repeat
16:   Apply Algorithm 1
17: until No more edges are removed from  $H$ 

```

1038 ► **Lemma 42.** *Suppose (H, Z) satisfies the invariants Invariants 1–3 and H has property*
 1039 *(Trim). If the execution of Algorithm 2 does not fail then the resulting H and Z also satisfy*
 1040 *the invariants and (Trim). If the execution fails then there is no H -deadlock scheme.*

1041 **Proof.** Suppose the algorithm finds a simple cycle $t_1 \xrightarrow{p_1} t_2 \cdots \xrightarrow{p_k} t_{k+1} = t_1$ where all p_i are
 1042 solid processes, and all t_i are distinct. By definition of a simple cycle, all p_i are distinct as
 1043 well. If there is a H -deadlock scheme then it should assign either $t_i \xrightarrow{p_i} t_{i+1}$ or $t_i \xleftarrow{p_i} t_{i+1}$ to
 1044 p_i .

1045 We examine the cases when the algorithm fails. The first reason for failure may appear
 1046 when all the edges on the cycle are blue. If for some j there is no reverse edge $t_j \xleftarrow{p_j} t_{j+1}$
 1047 in EH then a H -deadlock scheme, call it ds_H , should assign the edge $t_j \xrightarrow{p_j} t_{j+1}$ to p_j . In
 1048 consequence, as ds_H has to give each t_i at most one outgoing edge, all the edges in the cycle
 1049 should be in the image of ds_H . This is impossible as the cycle is blue. When there are reverse
 1050 edges $t_i \xleftarrow{p_i} t_{i+1} \in EH$ for all i , algorithm fails if all of them are blue. Indeed, there cannot
 1051 be an H -deadlock scheme in this case. The last reason for failure is when there is some solid
 1052 process p and p -labeled edges were removed by the algorithm. These must be edges of the
 1053 form $t_i \xrightarrow{p} t$ that are not on the cycle, for some $i = 1, \dots, k$. Those edges cannot be taken
 1054 in a deadlock scheme as it has to take the cycle in one direction or the other and thus cannot
 1055 take other edges from nodes on that cycle. As a deadlock scheme cannot assign any edge to
 1056 p , and p solid, there cannot be a deadlock scheme in that case.

1057 If the algorithm does not fail then either the cycle $t_1 \xrightarrow{p_1} t_2 \cdots \xrightarrow{p_k} t_{k+1} = t_1$ is not
 1058 blue, or its reverse is not blue. Let (H', Z') be the values after execution of the algorithm.
 1059 So $Z' = Z \cup \{t_1, \dots, t_k\}$, and H' is H after removing edges in line 11. We show that the
 1060 invariants hold.

1061 For Invariant 2, we observe that for every lock in Z' there is exactly one outgoing edge in
 1062 H' . So there is no solid edge from Z' to $H \setminus Z'$ if there was none from Z .

1063 For Invariant 3, we extend our Z -deadlock scheme to Z' . We choose the cycle found by
 1064 the algorithm or its reversal depending on which one is not blue. For every p_i we define
 1065 $ds_{Z'}(p_i)$ to be the edge in the chosen cycle. We set $ds_{Z'}(p) = \perp$ for all $p \in Proc_{Z'} \setminus Proc_Z$
 1066 other than p_1, \dots, p_k . We must show that such a p is necessarily fragile. Indeed, in this case
 1067 p must have one of its locks t in Z , and the other one, t' , in $Z' \setminus Z$. By Invariant 2, there is
 1068 no solid edge from t to t' in H . In H' all edges from t' to t are removed. So p is fragile as
 1069 the algorithm does not fail at line 12.

1070 For Invariant 1 suppose there is a deadlock scheme on H' . Then it is also a deadlock
 1071 scheme on H , as H' is a subgraph of H over the same set of locks. For the other direction take
 1072 a deadlock scheme ds_H on H . By Lemma 23, as we showed that Invariant 2 is maintained,
 1073 we can assume that ds_H is equal to $ds_{Z'}$ on Z' . We define a deadlock scheme $ds_{H'}$ on
 1074 H' . If $ds_H(p) = \perp$ then $ds_{H'}(p) = \perp$. If the source edge of $ds_H(p)$ is in $H \setminus Z'$ then
 1075 $ds_{H'}(p) = ds_H(p)$. This edge is guaranteed to exist in H' . If the two locks of p are both in
 1076 Z' let $ds_{H'}(p) = ds_H(p) = ds_{Z'}(p)$. The remaining case is when $ds(p)$ is an edge $t \xrightarrow{p} t'$ with
 1077 $t \in Z'$ and $t' \in H \setminus Z'$. In this case p is fragile as Z' has no solid arrows leaving it, and all
 1078 solid arrows in $Z' \setminus Z$ stay in Z' . We can then set $ds_{H'} = \perp$. It can be verified that $ds_{H'}$ is
 1079 a deadlock scheme. ◀

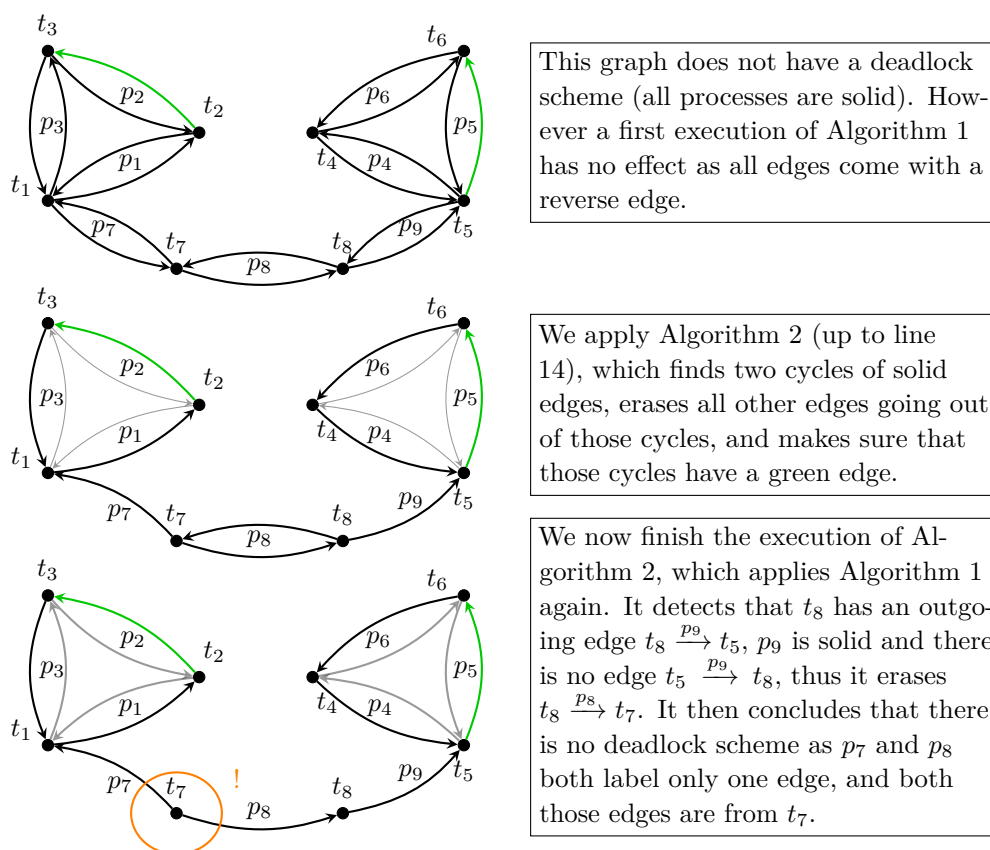
1080 ► **Lemma 43.** *If Algorithm 2 succeeds but does not increase Z or decrease H then (H, Z)
 1081 satisfies three properties:*

1082 **H1** H is trimmed.

1083 **H2** There is no cycle of solid edges intersecting $T \setminus Z$ in H .

1084 **H3** Every solid process has an edge in H .

1085 **Proof.** Since H was not modified, Algorithm 1 did not find any solo solid edge $t \xrightarrow{p} t'$ with
 1086 other outgoing edges from t , hence property H1 is satisfied.



■ **Figure 4** Illustration of Algorithm 1 and Algorithm 2.

1087 By Lemma 42, Invariant 2 is satisfied, hence any cycle intersecting $T \setminus Z$ in H must be
 1088 entirely in $T \setminus Z$. However if such a cycle existed then Algorithm 2 would not have stopped
 1089 on line 2 and thus would have either failed or increased Z . There are therefore no such cycle
 1090 intersecting $T \setminus Z$, hence property H2 is also satisfied.

1091 If H3 was not satisfied then Algorithm 2 would have failed on lines 12-13. ◀

1092 Since in the rest of the algorithm we increase Z and do not modify H , the three properties
 1093 form the lemma will continue to hold. We will refer to them as H1-H3.


1094 ▶ **Definition 44.** For a pair (H, Z) , we define an equivalence relation between locks in T :
 1095 $t_1 \equiv_H t_2$ if $t_1, t_2 \notin Z$ and there is a path of double solid edges between t_1 and t_2 .

1096 Intuitively, once we have trimmed the graph and eliminated simple cycles of solid edges
 1097 with Algorithm 2, the equivalence classes of \equiv_H are "trees" made of double solid edges.

1098 ▶ **Lemma 45.** If H satisfies property H1 and $t_1 \xrightarrow{p} t_2$ is in H for a solid process p then
 1099 either the \equiv_H -equivalence class of t_1 is a singleton, or $t_1 \xleftarrow{p} t_2$ is in H , hence $t_2 \equiv_H t_1$.

1100 **Proof.** Suppose there exists t_3 such that $t_1 \equiv_H t_3$, then there is a double solid edge from t_1 .
 1101 By (Trim) property, there cannot be solo outgoing edges from t_1 . ◀

1102 ▶ **Lemma 46.** Suppose H satisfies properties H1 and H2. Let $t_1, t_2 \in T \setminus Z$. If $t_1 \equiv_H t_2$
 1103 then there is a unique simple path of solid edges from t_1 to t_2 .



■ **Figure 5** Illustration of Algorithm 3, with dotted fragile edges and full solid ones. The black edges in the second figure are the ones selected in the deadlock scheme when we extend them in the proof of the algorithm.

1104 **Proof.** If $t_1 = t_2$ then any non-empty simple path of solid edges from t_1 to t_2 would contradict
 1105 property H2, hence the empty path is the only simple path from t_1 to t_2 .

1106 If $t_1 \neq t_2$ then by definition of \equiv_H there is a path of solid double edges from t_1 to t_2 ,
 1107 hence there is a simple path from t_1 to t_2 .

1108 Suppose there exist two distinct simple paths from t_1 to t_2 , then by Lemma 45 all the
 1109 locks on those paths are in the equivalence class of t_1 and t_2 . Hence as $t_1 \notin Z$, there is a
 1110 cycle of double solid edges intersecting $H \setminus Z$, contradicting property H2. ◀

1111 Our third algorithm looks for an edge $t_1 \xrightarrow{p} t_2$ with $t_1 \notin Z$ and $t_2 \in Z$, and adds the
 1112 full \equiv_H -equivalence class C of t_1 to Z . This step is correct, as we can extend a Z -deadlock
 1113 scheme to $(Z \cup C)$ -deadlock scheme by orienting edges in C .

■ **Algorithm 3** Extending Z with locks that can reach Z

```

1: while there exists  $t_1 \xrightarrow{p} t_2 \in HE$  with  $t_1 \notin Z$  and  $t_2 \in Z$  do
2:    $Z \leftarrow Z \cup \{t \in T \mid t \equiv_H t_1\}$ 
3: end while
    
```

1114 ▶ **Lemma 47.** Suppose H satisfies properties H1-H3, and (H, Z) satisfies Invariants 1–3.
 1115 After an execution of Algorithm 3 H and Z also have all these properties, and H has no
 1116 edges from $T \setminus Z$ to Z .

1117 **Proof.** Let (H', Z') be the pair obtained after execution of Algorithm 3. Observe that
 1118 $H' = H$, hence Invariant 1 holds. For the same reason H1 and H3 are still satisfied.
 1119 Furthermore, as Z can only increase, so H2 continues to hold.

1120 Let Z_m be the value of Z when entering the loop, and Z_m the value of Z at the end of
 1121 m -th iteration. So $Z_{m+1} = Z_m \cup \{t \in T \mid t \equiv_H t_1\}$, where $t_1 \xrightarrow{p} t_2$ is the edge found in the
 1122 guard of the while statement. We verify that Z_{m+1} satisfies Invariants 2 and 3 if Z_m does.

1123 For Invariant 2, Lemma 45 says that there are no outgoing solid edges from the \equiv_H -
 1124 equivalence class of t_1 , unless that class is a singleton. If it is a singleton, there are no
 1125 outgoing solid edges from t_1 or $t_1 \xrightarrow{p} t_2$ is the only outgoing edge of t_1 . In both cases, there
 1126 are no solid edges from Z_{m+1} to $T \setminus Z_{m+1}$ in H .

1127 For Invariant 3 we extend a Z_m -deadlock scheme to Z_{m+1} . So we are given ds_m and
 1128 construct ds_{m+1} . If the two locks of some process q are in Z then $ds_{m+1}(q) = ds_m(q)$. We
 1129 set $ds_{m+1}(q)$ to be the edge $t_1 \xrightarrow{q} t_2$ found by the algorithm. Let C be the equivalence class
 1130 of t_1 : $C = \{t \in T \mid t \equiv_H t_1\}$. By Lemma 46 there is a unique simple path from $t \in C$ to
 1131 t_1 . Let $t \xrightarrow{q} t'$ be the first edge on this path. We set $ds_{m+1}(q)$ to be this edge. We set
 1132 $ds_{m+1}(q) = \perp$ for all remaining processes q .

1133 We verify that ds_{m+1} is a Z_{m+1} -deadlock scheme. By the above definition every lock in
 1134 C has a unique outgoing edge in ds_{m+1} , hence every lock in Z_{m+1} does. It is also immediate
 1135 that the image of ds_{m+1} does not contain a blue cycle as it would need to be already in the
 1136 image of ds_m (every lock has exactly one outgoing edge in ds_{m+1} and the path obtained
 1137 by following those edges from an element of C leads to Z_m). It is maybe less clear that
 1138 $ds_{m+1} \neq \perp$ for every solid $q \in Proc_{Z_{m+1}}$. Let q be a solid process in $Proc_{Z_{m+1}}$, and suppose
 1139 ds_{m+1} is not defined by the procedure from the previous paragraph. If both of locks of q
 1140 are in Z_m then $ds_{m+1}(q)$ must be defined because $ds_m(q)$ is. If $q = p$, the process labeling
 1141 the transition chosen by the algorithm, then $ds_{m+1}(q)$ is defined. Otherwise both locks of q
 1142 are in C . Say these are t and t' . If neither $t \xrightarrow{q} t'$ is on the shortest path from t to t_1 , nor
 1143 is $t \xrightarrow{q} t'$ on the shortest path from t' to t_1 then there must be a cycle in C . But this is
 1144 impossible as we assumed that there are no cycles of solid edges intersecting $T \setminus Z$ (property
 1145 H2) and $Z \subseteq Z_m$. Hence $ds_{m+1}(q)$ is defined, and ds_{m+1} is a Z_{m+1} -deadlock scheme.

1146 All that is left to prove is that H has no edges from $T \setminus Z$ to Z , which is immediate as
 1147 otherwise Algorithm 3 would not have stopped. ◀

■ Algorithm 4 Incorporating green cycles

1: **if** there exists a green cycle $t_1 \xrightarrow{p_1} t_2 \cdots \xrightarrow{p_k} t_{k+1} = t_1$ with $t_1 \xrightarrow{p_k} t_1$ green **then**
 2: $Z \leftarrow Z \cup \bigcup_{i=1}^k \{t \mid t \equiv_H t_i\}$
 3: **end if**

1148 ► **Lemma 48.** *Suppose H satisfies H1-H3, (H, Z) satisfies Invariants 1–3, and moreover*
 1149 *there are no edges from $T \setminus Z$ to Z . After an execution of Algorithm 4 H satisfies H1-H3,*
 1150 *and (H, Z) satisfies Invariants 1–3.*

1151 **Proof.** Let (H', Z') be the pair obtained after execution of Algorithm 3. Observe that
 1152 $H' = H$, hence Invariant 1 holds. For the same reason H1 and H3 are still satisfied.
 1153 Furthermore, as Z can only increase, so is H2. It remains to verify Invariants 2 and 3.

1154 Consider the green cycle found by the algorithm: $t_1 \xrightarrow{p_1} t_2 \cdots \xrightarrow{p_k} t_{k+1} = t_1$

1155 For Invariant 2, Lemma 45 says that there are no outgoing solid edges from the \equiv_H -
 1156 equivalence class of t_1 , unless that class is a singleton. If it is a singleton, there are no
 1157 outgoing solid edges from t_1 or $t_1 \xrightarrow{p} t_2$ is the only outgoing edge of t_1 . In both cases, there
 1158 are no solid edges from Z_{m+1} to $T \setminus Z_{m+1}$ in H .

1159 For Invariant 3 we extend a Z -deadlock scheme ds_Z to Z' . For every lock $t \in Z' \setminus Z$ let
 1160 j be the biggest index among $1, \dots, k$ with $t \equiv_H t_j$. If $t = t_j$ then set $ds_{Z'}(p_j)$ to be the

1161 edge $t_j \xrightarrow{p_j} t_{j+1}$. Otherwise, take the unique path from t to t_j in the \equiv_H -equivalence class of
 1162 the two locks; this is possible thanks to Lemma 46. If the path starts with $t \xrightarrow{p} t'$ then set
 1163 $ds_{Z'}(p)$ to this edge. Then set $ds_{Z'}(p) = \perp$ for all remaining processes p .

1164 We claim that $ds_{Z'}$ is a Z' -deadlock scheme. First, there is an outgoing $ds_{Z'}$ edge from
 1165 every lock in Z' because of the definition. Moreover it is unique.

1166 We need to show that $ds_{Z'}(p)$ is defined for every solid process p . This is clear if the two
 1167 locks, t and t' , of p are in Z . If both locks are not in Z then either $t \equiv_H t'$ or there is a solo
 1168 solid edge between the two, say $t \xrightarrow{p} t'$. In the latter case this is the only edge from t , as H is
 1169 trimmed. As the equivalence class of t is then a singleton, this must be an edge on the cycle
 1170 and $ds_{Z'}(p)$ is defined to be this edge. Suppose $t \equiv_H t'$ and $ds_{Z'}(p)$ is not defined. Let j be
 1171 the biggest index among $1, \dots, k$ such that $t \equiv_H t_j$. If neither $t \xrightarrow{p} t'$ is on the shortest path
 1172 from t to t_j , nor $t \xleftarrow{p} t'$ is on the shortest path from t' to t_j then there must be a cycle in
 1173 C . But this is impossible as we assumed that there are no cycles of solid edges intersecting
 1174 $T \setminus Z_m$ in H (property H2). The remaining case is when one of the locks of p is in Z and
 1175 another in $Z' \setminus Z$. There is no solid edge leaving Z by Invariant 2. There is no solid edge
 1176 entering Z by the assumption of the lemma. So p is a solid process labeling no edge in H
 1177 which contradicts H3.

1178 The last thing to verify is that there is no blue cycle in $ds_{Z'}$. We first check that $ds_{Z'}$
 1179 contains $t_k \xrightarrow{p_k} t_1$. This is because t_k is necessary the last from its equivalence class. A
 1180 blue cycle cannot contain locks from Z as there are no edges entering Z in $ds_{Z'}$. Let
 1181 $t'_1 \xrightarrow{p'_1} t'_2 \dots \xrightarrow{p'_l} t'_{l+1} = t'_1$ be a hypothetical blue cycle in $Z' \setminus Z$ using transitions in $ds_{Z'}$.

1182 Consider x such that $t'_1 \equiv_H t'_j$ for $j \leq x$ but $t'_1 \not\equiv_H t'_{x+1}$. By definition of $ds_{Z'}$ we must
 1183 have that t'_x is the last lock among t_1, \dots, t_k equivalent to t'_1 , say it is t_y . As each lock only
 1184 has one outgoing transition in the image of $ds_{Z'}$, and as there is a path from t_y to t_k in that
 1185 image, t_k must be on that cycle, and thus the green edge $t_k \xrightarrow{p_k} t_1$ as well, contradicting the
 1186 assumption that this is a blue cycle. ◀

1187 Algorithm 5 below is the complete algorithm as required by Proposition 24. Its correctness
 1188 is stated in the next lemma.

■ **Algorithm 5** Complete algorithm

```

1:  $H \leftarrow G_\sigma$ 
2:  $Z \leftarrow \emptyset$ 
3: repeat
4:   Apply Algorithm 1
5: until No more edges are removed from  $H$ 
6: repeat ▷  $H$  is trimmed
7:   Apply Algorithm 2
8: until No more edges are removed from  $H$ 
9: repeat ▷ From now on  $H$  satisfies properties H1-H3
10:  Apply Algorithm 3 ▷ no edges from  $T \setminus Z$  to  $Z$ 
11:  Apply Algorithm 4
12: until  $Z$  does not increase anymore
13: if there is a process  $p \in Proc \setminus Proc_Z$  that is not  $Z$ -lockable then
14:   return "Fail:  $\sigma$  is winning"
15: else
16:   return " $\sigma$  is not winning"
17: end if

```

1189 ► **Lemma 49.** *Algorithm 5 fails if and only if $(G_\sigma, \{L_p\}_{p \in Proc})$ admits a deadlock scheme.*
 1190 *The algorithm works in polynomial time.*

1191 **Proof.** Let (H', Z') be the values at the end of the execution of the algorithm.

1192 Suppose the algorithm fails. If it is before line 13 then using the previous lemmas and
 1193 Invariant 1 we get that G_σ does not have a deadlock scheme. If the algorithm fails in line 14
 1194 then there exists a process p with one of its locks outside of Z and not Z -lockable. Suppose
 1195 towards a contradiction that there is a deadlock scheme ds_H for H . It must have an edge
 1196 from a lock of p that is not in Z , say from t . By definition, every lock with an incoming
 1197 edge in ds_H must also have an outgoing edge in ds_H . Following these edges we get a cycle in
 1198 H . During the last iteration of lines 9-12, Z was not increased, hence by Lemma 47 there
 1199 are no edges from $T \setminus Z$ to Z . This cycle is therefore outside Z . It has to be a green cycle
 1200 by definition of a deadlock scheme, which is a contradiction because Algorithm 4 did not
 1201 increase Z in its last application.

1202 If the algorithm succeeds then there is a Z -deadlock scheme, say ds_Z . It may still not
 1203 be a deadlock scheme on G_σ because $Proc_Z$ may be a strict subset of $Proc$. We construct a
 1204 deadlock scheme (Z, ds) for G_σ . First, we set $ds(p) = ds_Z(p)$ for all $p \in Proc$. Let us take
 1205 $p \in Proc \setminus Proc_Z$, as the algorithm did not fail in lines 13-14, p is Z -lockable and thus we
 1206 can map it to \perp .

1207 Finally, this algorithm operates in polynomial time as all steps of all loops in the algorithms
 1208 either decrease H or increase Z . Furthermore, the condition on line 13 is easily verifiable
 1209 by checking in the lockset family $\{L_p\}_{p \in Proc}$ of σ whether there exists $B \in L_p$ such that
 1210 $B \subseteq Z$. ◀

1211 **C Solving the exclusive case in PTIME**

1212 ► **Lemma 25.** *Let σ be a locally live control strategy and G_σ its lock graph. For all $t_1, t_2 \in T$,*
 1213 *if G_σ has a blue edge $t_1 \xrightarrow{p} t_2$ then it has a green edge $t_2 \xrightarrow{p} t_1$.*

1214 **Proof.** Suppose there is a blue edge $t_1 \xrightarrow{p} t_2$, then there is a process p and a local run of \mathcal{A}_p
 1215 of the form $u = u_1(a_1, \mathbf{acq}_{t_1})u_2(a_1, \mathbf{rel}_{t_2})u_3(a_3, \mathbf{acq}_{t_2})$ respecting σ , with no \mathbf{rel}_{t_1} in u_2 or
 1216 u_3 . Hence there is a point in the run at which p holds both locks.

1217 It is not possible that there is always a release between two acquire operations in u , as
 1218 then the run would never hold both locks. So there are two acquires in u with no release
 1219 in-between. Thus there is a green edge between the first lock taken and the second one
 1220 because 2LSS is exclusive. As $t_1 \xrightarrow{p} t_2$ is blue, the only possibility is that there is a green
 1221 edge $t_2 \xrightarrow{p} t_1$. ◀

1222 ► **Lemma 26.** *Let σ be a locally live control strategy and G_σ its lock graph. For every edge*
 1223 *$t_1 \xrightarrow{p} t_2$ in G , process p is $\{t_1, t_2\}$ -lockable.*

1224 **Proof.** By definition of G_σ , as $t_1 \xrightarrow{p} t_2$ is an edge, there exists a local σ -run u_p of p acquiring
 1225 t_1 . Consider the first acquire transition in the run $u_p = u(a, \mathbf{acq}_{t_i})u'$ for some $i \in \{1, 2\}$ and
 1226 u containing only local actions. As our 2LSS is exclusive, this means that u makes p reach
 1227 a configuration with only one outgoing transition acquiring t_i , and p not having any lock.
 1228 Hence p is $\{t_i\}$ -lockable and thus $\{t_1, t_2\}$ -lockable. ◀

1229 ► **Proposition 27.** *A strategy σ is winning if and only if there exists a process that is not*
 1230 *BT_σ -lockable.*

1231 The left-to-right direction is handled by the lemma below.

1232 ► **Lemma 50.** *If all processes are BT_σ -lockable then σ is not winning.*

1233 **Proof.** We construct a BT_σ -deadlock scheme for $(G_\sigma, Locks_\sigma)$ as follows: for all $t \in BT_\sigma$ we
 1234 select an outgoing edge in BT_σ , say $t \xrightarrow{P_t} \bar{t}$ to some $\bar{t} \in BT_\sigma$, and a green one if possible.
 1235 Such an edge always exists by definition of BT_σ . We define ds as $ds(p_t) = t \xrightarrow{P_t} \bar{t}$ for all
 1236 $t \in BT_\sigma$, and $ds(p) = \perp$ for all other $p \in Proc$.

1237 We show that ds is a BT_σ -deadlock scheme for $(G_\sigma, Locks_\sigma)$. Hence it is also just a
 1238 deadlock scheme because BT_σ locks all processes.

1239 Clearly for all $p \in Proc$, $ds(p)$ is either \perp or a p -labeled edge within BT_σ . Furthermore as
 1240 all processes are BT_σ -lockable, in particular the ones mapped to \perp by ds are. It is also clear
 1241 that all locks of BT_σ have an unique outgoing edge. Now suppose there is a blue cycle in
 1242 $ds(Proc)$, then by Lemma 25 there is a reverse cycle of green edges between the same locks.
 1243 This means all those locks have an outgoing green edge within BT_σ , which is a contradiction
 1244 as we have chosen for ds green outgoing edges whenever possible. ◀

1245 For the right-to-left direction we first prove an auxiliary lemma.

1246 ► **Lemma 51.** *If ds is a Z -deadlock scheme for $(G_\sigma, Locks_\sigma)$ then $Z \subseteq BT_\sigma$.*

1247 **Proof.** Suppose there exists $t \in Z \setminus BT_\sigma$, then there exists p such that $ds(p) = t \xrightarrow{P} t'$
 1248 for some t' . By definition of BT_σ , there are no edges from $T \setminus BT_\sigma$ to BT_σ in G_σ , hence
 1249 $t' \in Z \setminus BT_\sigma$. By iterating this process we eventually discover a proper cycle in G_σ outside
 1250 of BT_σ , which is impossible as this cycle should be part of a direct deadlock component, and
 1251 thus be included in BT_σ . ◀

1252 ► **Lemma 52.** *If some process p is not BT_σ -lockable then σ is winning.*

1253 **Proof.** Suppose there exists p that is not BT_σ -lockable.

1254 Towards a contradiction assume that σ is not winning, hence there is a Z -deadlock scheme
 1255 ds for σ . As p is not BT_σ -lockable, it is not Z -lockable either, hence $ds(p)$ is not \perp . So $ds(p)$
 1256 must be an edge $t_1 \xrightarrow{P} t_2$ from G_σ and $t_1, t_2 \in Z$. By previous lemma $t_1, t_2 \in BT_\sigma$.

1257 By Lemma 26, p is $\{t_1, t_2\}$ -lockable, and therefore also BT_σ -lockable, yielding a contra-
 1258 diction. ◀

1259 This completes the proof of Proposition 27.

1260 ► **Proposition 28.** *Let \mathcal{S} be an exclusive 2LSS. There is a winning locally live strategy for
 1261 the system if and only if there exists a locally live strategy σ_p for some process p preventing
 1262 it from acquiring any lock from $BT_{\mathcal{S}}$.*

1263 **Proof.** One direction is easy: if all strategies make all processes able to acquire a lock from
 1264 $BT_{\mathcal{S}}$ then there is no winning strategy. Let σ be a control strategy, and G_σ its lock graph
 1265 and its SCCs. Note that $G_{\mathcal{S}}$ is a subgraph of G_σ , hence every SCC in G_σ is a superset of
 1266 an SCC in $G_{\mathcal{S}}$. Observe that if an SCC in G_σ contains a direct semi-deadlock SCC of $G_{\mathcal{S}}$
 1267 then it is direct deadlock. Indeed, if an SCC in $G_{\mathcal{S}}$ is a direct semi-deadlock but not a direct
 1268 deadlock then σ adds an edge $t_1 \xrightarrow{P} t_2$ to this SCC in G_σ . As t_1, t_2 are in that SCC of $G_{\mathcal{S}}$,
 1269 there is a simple path from t_2 to t_1 not involving p . Hence, a direct semi-deadlock SCC
 1270 becomes a direct-deadlock SCC. This implies $BT_{\mathcal{S}} \subseteq BT_\sigma$. Let $p \in Proc$, as there is a run of
 1271 p acquiring a lock of $BT_{\mathcal{S}}$, either p is $BT_{\mathcal{S}}$ -lockable (and thus BT_σ -lockable) or there is an
 1272 edge labeled by p towards $BT_{\mathcal{S}}$, meaning that both locks of p are in BT_σ and thus that p is
 1273 BT_σ -lockable by Lemma 26. As a consequence, all processes are BT_σ -lockable. We conclude
 1274 by Proposition 27.

1275 In the other direction we suppose that there exists a process p and a strategy σ_p forbidding
 1276 p from acquiring a lock of BT_S . We construct a strategy σ such that p is not BT_σ -lockable.
 1277 This will show that σ is winning.

1278 Let $FT = T \setminus BT_S$ be the set of locks not in BT_S . In G_S , no node of FT can reach
 1279 a direct semi-deadlock SCC. In particular, there is no direct semi-deadlock SCC in G_S
 1280 restricted to FT . We construct a strategy σ such that, when restricted to FT , the SCC's of
 1281 G_σ and G_S are the same.

1282 Let us linearly order SCC components of G_S restricted to FT in such a way that if a
 1283 component C_1 can reach a component C_2 then C_1 is before C_2 in the order.

1284 We use strategy σ_p for p . For every process $q \neq p$ we have one of the two cases: (i) either
 1285 there is a local strategy σ_q inducing only the edges that are already in G_S ; (ii) or every local
 1286 strategy induces some edge that is not in G_S . In the second case there are no q -labeled edges
 1287 in G_S , and for each of the two possible edges there is a strategy inducing only this edge.

1288 For a process q from the first case we take a strategy σ_p that induces only the edges
 1289 present in G_S .

1290 For a process q from the second case,

- 1291 ■ If both locks of q are in BT_S then take any strategy for p .
- 1292 ■ If one of the locks of q is in BT_S and the other in FT then choose a strategy inducing an
 1293 arrow from the BT_S lock to the FT lock.
- 1294 ■ If both locks of q are in FT then choose a strategy inducing an edge from a smaller to a
 1295 bigger SCC.

1296 Consider the graph G_σ of the resulting strategy σ . Restricted to FT this graph has the
 1297 same SCCs as G_S . Moreover, there are no extra edges in G_σ added to any SCC included in
 1298 FT , and there are no edges from FT to BT_S . As a result, we have $BT_S = BT_\sigma$.

1299 As p cannot acquire any lock from BT_S , it is not BT_S -lockable and thus not BT_σ -lockable
 1300 either. ◀

1301 **D** Nested-locking strategies

1302 By abuse of language we will denote in this section a run u , not necessarily initial, as neutral
 1303 if every lock acquired in u is also released within u .

1304 ► **Lemma 53.** *Every local run u that respects a nested-locking strategy has a unique stair*
 1305 *decomposition.*

1306 **Proof.** We set $u = u_1 \text{acq}_{t_1} u_2 \text{acq}_{t_2} \cdots u_k \text{acq}_{t_k} u_{k+1}$ such that $\{t_1, \dots, t_k\}$ is the set of locks
 1307 held by p at the end of the run, and the distinguished acq_{t_i} are the last time these locks
 1308 were taken in u . Consequently, there is no operation on t_i in u_{i+1}, \dots, u_{k+1} .

1309 Observe that u_{k+1} must be neutral because the process owns $\{t_1, \dots, t_k\}$ at the end
 1310 of u . If some u_i , $i \leq k$, is not neutral, then there exists $t \in T$ such that $t \notin \{t_1, \dots, t_i\}$
 1311 and p holds t after $u_1 \text{acq}_{t_1} \cdots u_i \text{acq}_{t_i}$. Then p has to release t at some point later in the
 1312 run: if $t \notin \{t_1, \dots, t_k\}$ then p does not hold it at the end, and if $t \in \{t_1, \dots, t_k\}$ then
 1313 $t \in \{t_{i+1}, \dots, t_k\}$ and thus t is taken again later in the run. But this contradicts the nested-
 1314 locking assumption, because t would be released before t_i , which has been acquired after
 1315 t .

1316 ◀

1317 ► **Lemma 31.** *A nested-locking control strategy σ with behavior $(\mathcal{P}_p)_{p \in Proc}$ is **not** winning*
 1318 *if and only if for every $p \in Proc$ there is a stair pattern $(T_{owns}^p, T_{blocks}^p, \preceq^p) \in \mathcal{P}_p$ such that:*

- 1319 ■ $\bigcup_{p \in Proc} T_{locks}^p \subseteq \bigcup_{p \in Proc} T_{owns}^p$,
 1320 ■ the sets T_{owns}^p are pairwise disjoint,
 1321 ■ there exists a total order \preceq , on the set of all locks T , compatible with all \preceq^p .

1322 **Proof.** Suppose σ is not winning, let w be a run leading to a deadlock. For all p let T_{owns}^p
 1323 be the set of locks owned by p after w . Take u^p the local run of p in w . Since w leads to a
 1324 deadlock every u^p is risky. For every p , consider the stair pattern $(T_{owns}^p, T_{locks}^p, \preceq^p)$ of u^p .
 1325 This way we ensure it is a pattern from \mathcal{P}_p .

1326 We need to show that these patterns satisfy the requirements of the lemma. Since the
 1327 configuration reached after w is a deadlock, every process waits for locks that are already
 1328 taken so $T_{locks}^p \subseteq \bigcup_{q \in Proc} T_{owns}^q$, for every process p , proving the first condition.

1329 We have that T_{owns}^p is the set of locks that p has at the end of the run w . So the sets
 1330 T_{owns}^p are pairwise disjoint.

1331 For the last requirement of the lemma take an order \preceq on T satisfying: $t \preceq t'$ if the last
 1332 operation on t appears before the last operation on t' in w .

1333 Let $p \in Proc$, let $u^p = u_1^p \mathbf{acq}_{t_1^p} u_2^p \mathbf{acq}_{t_2^p} \cdots u_k^p \mathbf{acq}_{t_k^p} u_{k+1}^p$ be the stair decomposition of u^p .
 1334 As p never releases t_i^p , the distinguished $\mathbf{acq}_{t_i^p}$, is the last operation on t_i^p in the global run.
 1335 Consequently, for all t we have $t_i^p \preceq t$ whenever t is used in $u_{i+1}^p \mathbf{acq}_{t_{i+1}^p} \cdots u_k^p \mathbf{acq}_{t_k^p} u_{k+1}^p$. As
 1336 a result, \preceq is compatible with all \preceq^p .

1337 For the converse implication, suppose that there are patterns satisfying all the conditions
 1338 of the lemma. We need to construct a run w ending in a deadlock. For every process
 1339 p we have a stair pattern $(T_{owns}^p, T_{locks}^p, \preceq^p)$ coming from a local σ -run u^p of p , with
 1340 $u = u_1^p \mathbf{acq}_{t_1^p} u_2^p \mathbf{acq}_{t_2^p} \cdots u_k^p \mathbf{acq}_{t_k^p} u_{k+1}^p$ as stair decomposition. There is also a linear order \preceq
 1341 compatible with all \preceq_p . Let \prec be its strict part. Let $t_1 \dots t_k$ be the sequence of locks from
 1342 $\bigcup_p T_p$ listed in \prec order (which is possible as the T_{owns} are disjoint and thus no lock appears
 1343 twice in that sequence), and let $\{p_1, \dots, p_n\} = Proc$. We claim that we can get a suitable run
 1344 w as $u_1^{p_1} \dots u_n^{p_n} w'$ where w' is obtained from $t_1 \dots t_k$ by substituting each t_i^p by $\mathbf{acq}_{t_i^p} u_{i+1}^p$.

1345 All u_i^p are neutral, hence after executing $u_1^{p_1} \dots u_n^{p_n}$ all locks are free. Let $t_i^p \in T_p$,
 1346 suppose we furthermore executed all $\mathbf{acq}_{t_j^q} u_{j+1}^q$ with $t_j^q \prec t_i^p$. Then the set of non-free locks
 1347 is $\{t_j^q \mid t_j^q \prec t_i^p\}$. As \preceq is compatible with all \preceq^p , all locks t used in $\mathbf{acq}_{t_i^p} u_{i+1}^p$ are such that
 1348 $t_i^p \preceq t$. Moreover, all t_j^q that were taken before are such that $t_j^q \prec t_i^p$, thus $\mathbf{acq}_{t_i^p} u_{i+1}^p$ only
 1349 uses locks that are free and can therefore be executed.

1350 As a result, w can be executed. It leads to a deadlock as $T_{locks}^p \subseteq \bigcup_q T_{owns}^q$. ◀

1351 ► **Lemma 32.** Given a lock-sharing system $((\mathcal{A}_p)_{p \in Proc}, \Sigma^s, \Sigma^e, T)$, a process $p \in Proc$ and
 1352 a set of patterns \mathcal{P}_p , we can check in polynomial time in $|\mathcal{A}_p|$ and $2^{|T|}$ whether there exists a
 1353 nested-locking local strategy σ_p with set of patterns included in \mathcal{P}_p .

1354 **Proof.** For every process p we proceed as follows We extend the states of p to store in it the
 1355 stair profile of the current run. This increases the number of states by the factor $|T|!2^{|T|^2}$.
 1356 As the set of locks owned by p is now a function of the current state, this also allows us to
 1357 eliminate all non-realizable transitions which acquire a lock that p owns or release one it
 1358 does not have.

1359 Then the risky nature and the stair pattern of a run u depend only on the state it ends
 1360 in and the choice of transitions of the system from that state after executing u .

1361 We focus on states with only outgoing transitions acquiring locks (including those with
 1362 no outgoing transition). Those states force the runs entering them to be risky. The choices
 1363 of transitions of the system in that state will then determine the stair pattern of each run

1364 entering it. We mark such states as bad if all choices of transitions of the system yield a
1365 stair pattern that is not in \mathcal{P}_p .

1366 We iteratively delete all bad states and all their ingoing transitions, as we need to ensure
1367 that we never reach them. If we delete an uncontrollable transition then we mark its source
1368 state as bad as reaching that state would make the environment able to reach a bad state.
1369 As we deleted some transitions, we may have reduced the choices of the system in some
1370 states, therefore we check again all states and mark as bad all the ones in which no choice of
1371 transitions of the system yields a pattern in \mathcal{P}_p . We again delete all bad states, and so on.

1372 At some point we reach a system with no bad state. If it is empty, then we conclude that
1373 there is no suitable strategy. If not, we construct a strategy by selecting for each state a set
1374 of outgoing transitions that either contains a transition not acquiring a lock or allows runs
1375 reaching this state to match a pattern of \mathcal{P}_p . This strategy ensures that all risky runs have a
1376 pattern in \mathcal{P}_p . ◀

1377 ▶ **Proposition 54.** *The deadlock avoidance control problem is NEXPTIME-hard.*

1378 **Proof.** For the lower bound, we reduce the domino tiling problem over an exponential grid.
1379 In this problem, we are given an alphabet Σ , with a special letter b , an integer n (in unary)
1380 and a set D of dominoes, each domino d being a 4-tuple $(up_d, down_d, right_d, left_d)$ of letters
1381 of Σ . The output is whether there exists a function $til : \{0, \dots, 2^n - 1\}^2 \rightarrow D$ such that for
1382 all $x, y, x', y' \in \{0, \dots, 2^n - 1\}$,

- 1383 ■ if $x' = x$ and $y' = y + 1$ then $up_{til(x,y)} = down_{til(x',y')}$,
- 1384 ■ if $x' = x + 1$ and $y' = y$ then $right_{til(x,y)} = left_{til(x',y')}$,
- 1385 ■ if $x = 0$ then $left_{til(x,y)} = b$
- 1386 ■ if $x = 2^n - 1$ then $right_{til(x,y)} = b$
- 1387 ■ if $y = 0$ then $down_{til(x,y)} = b$
- 1388 ■ if $y = 2^n - 1$ then $up_{til(x,y)} = b$

1389 This problem is well-known to be NEXPTIME-complete.

1390 Let n, Σ, D, b be an instance of that problem. We construct a lock-sharing system as
1391 follows: We have three processes p, \bar{p} and q . Process p will use locks $0_i^x, 1_i^x, 0_i^y, 1_i^y$, for
1392 $1 \leq i \leq n$, together with $t(d)$ for each domino $d \in D$, and an extra lock $lock$. Process p will
1393 use analogous locks but with a bar: $\bar{0}_i^x, \bar{1}_i^x, \bar{0}_i^y, \bar{1}_i^y, \bar{t}(d), \bar{lock}$. Process q will use all the locks.

1394 Let us describe process q . In the initial state the environment can choose between several
1395 actions: *equality*, *vertical*, *horizontal*, b_{left} , b_{right} , b_{up} and b_{down} . Each of these actions
1396 leads to a different transition system, but the principle behind all systems is the same. In
1397 the first phase, for each $1 \leq i \leq n$, the environment can choose to take either 0_i^x or 1_i^x , and
1398 then take either $\bar{0}_i^x$ or $\bar{1}_i^x$. In the second phase the same happens for y locks. After these two
1399 phases environment has chosen two pairs of n -bit numbers, call them $\#x, \#y$ and $\#\bar{x}, \#\bar{y}$.
1400 Where the three systems differ is how the choice of \bar{x} 's and \bar{y} 's is limited in these two phases.
1401 This depends on the first action done by the environment.

- 1402 ■ If it is *equality* then $\#x = \#\bar{x}$ and $\#y = \#\bar{y}$.
- 1403 ■ If it is *vertical*, then $\#x = \#\bar{x}$ and $\#y + 1 = \#\bar{y}$;
- 1404 ■ If it is *horizontal*, then $\#x + 1 = \#\bar{x}$ and $\#y = \#\bar{y}$.
- 1405 ■ If it is b_{left} (resp. b_{right}) then $\#x = 0$ (resp. $2^n - 1$).
- 1406 ■ If it is b_{down} (resp. b_{up}) then $\#y = 0$ (resp. $2^n - 1$).

1407 These constraints are easily implemented. For example, *equality* is checked by forcing the
1408 environment to take the same bit for \bar{x} after choosing each bit for x (similarly for y).

1409 In the third phase, process q has to take and then immediately release $lock$ and then \bar{lock}
1410 before it reaches a state *dominoes*.

1411 Every state in the three phases before *dominoes* has a local loop on it, meaning that q
 1412 cannot deadlock while being in one of these states.

1413 In state *dominoes*, the system chooses to take two dominoes d and \bar{d} such that:

- 1414 ■ If the environment chose *equality* then $d = \bar{d}$
- 1415 ■ If it chose *vertical* then $up_d = down_{\bar{d}}$
- 1416 ■ If it chose *horizontal* then $right_d = left_{\bar{d}}$
- 1417 ■ If it chose b_{left} (resp. $b_{right}, b_{up}, b_{down}$) then $left_d = b$ (resp. $right_d, up_d, down_d$).

1418 Each choice leads to a different state $s_{d,\bar{d}}$. From there transitions force the system to
 1419 take every lock $t(d')$ and $\bar{t}(d')$, except for $t(d)$ and $\bar{t}(d)$ in order to reach a state *win* with a
 1420 local loop on it and no other outgoing transitions.

1421 We now describe process p . It starts by taking the lock *lock*, which it never releases.
 1422 Then the environment chooses to take one of 0_i^x and 1_i^x and one of 0_i^y and 1_i^y for all $1 \leq i \leq n$.
 1423 Finally, the system chooses a domino d and takes the lock $t(d)$ before reaching a state with
 1424 no outgoing transitions.

1425 Process \bar{p} behaves identically, but uses locks with a bar.

1426 We need to show that if there is a tiling $til : \{0, \dots, 2^n - 1\}^2 \rightarrow D$ then there is a winning
 1427 strategy. The strategy for q is to respond with the correct tiles: if environment chooses $\#x$,
 1428 $\#y$, $\#\bar{x}$, $\#\bar{y}$ the strategy chooses locks corresponding to d_1 and \bar{d}_2 with $d_1 = til(\#x, \#y)$ and
 1429 $d_2 = til(\#\bar{x}, \#\bar{y})$. The strategy of p does the same but uses inverse encoding of numbers:
 1430 considers 0 as 1, and 1 as 0. Similarly for \bar{p} .

1431 Assume for contradiction that the strategy is not winning, so we have a run leading to a
 1432 deadlock. First, observe that the environment must have q pass the *lock* phase before p and
 1433 \bar{p} start running, because all states before *lock* have a self-loop, so q cannot block there. If p
 1434 or \bar{p} starts before q has passed the *lock* phase, then q can never pass it as one of $lock, \bar{lock}$
 1435 will never be available again.

1436 If q passed the *lock* phase then process p has no choice but to take *lock*, and then the
 1437 remaining locks among x, y . Similarly for \bar{p} . At this stage strategy σ is defined so that the
 1438 three processes will never take the same lock. So q cannot be blocked before reaching state
 1439 *win*. Thus deadlock is impossible.

1440 For the other direction, suppose there is a winning strategy σ for the system. Observe
 1441 that the strategy σ_p for process p should decide which domino to take after the environment
 1442 has decided what x and y locks to take. So σ_p defines a function $til : \{0, \dots, 2^n - 1\}^2 \rightarrow D$.
 1443 Similarly $\sigma_{\bar{p}}$ defines \bar{til} .

1444 We first show that $til(i, j) = \bar{til}(i, j)$ for all $i, j \in \{0, \dots, 2^n - 1\}$. If not then consider the
 1445 run where environment chooses *equality* and then x, \bar{x} to be the representations of i , and
 1446 y, \bar{y} to be representations of j . So q reaches state *dominoes*. Next the environment makes
 1447 processes p and \bar{p} to choose locks corresponding to dominoes $til(i, j)$ and $\bar{til}(i, j)$. The two
 1448 processes p and q reach a deadlock state. Since these are two different dominoes, q cannot
 1449 reach state *win* from any state $s_{d,\bar{d}}$. Hence there is a deadlock run, that we have assumed
 1450 impossible.

1451 Once we know that the strategies σ_p and $\sigma_{\bar{p}}$ define the same tiling function it is easy
 1452 to see that in order to be winning when environment chooses *vertical*, *horizontal* or b_{left} ,
 1453 b_{right} , b_{down} , b_{up} actions, the tiling function should be correct.

1454 ◀




Figure 6 Transition system for process p (with $D = \{d_1, \dots, d_m\}$), dashed arrows are controlled by the system.

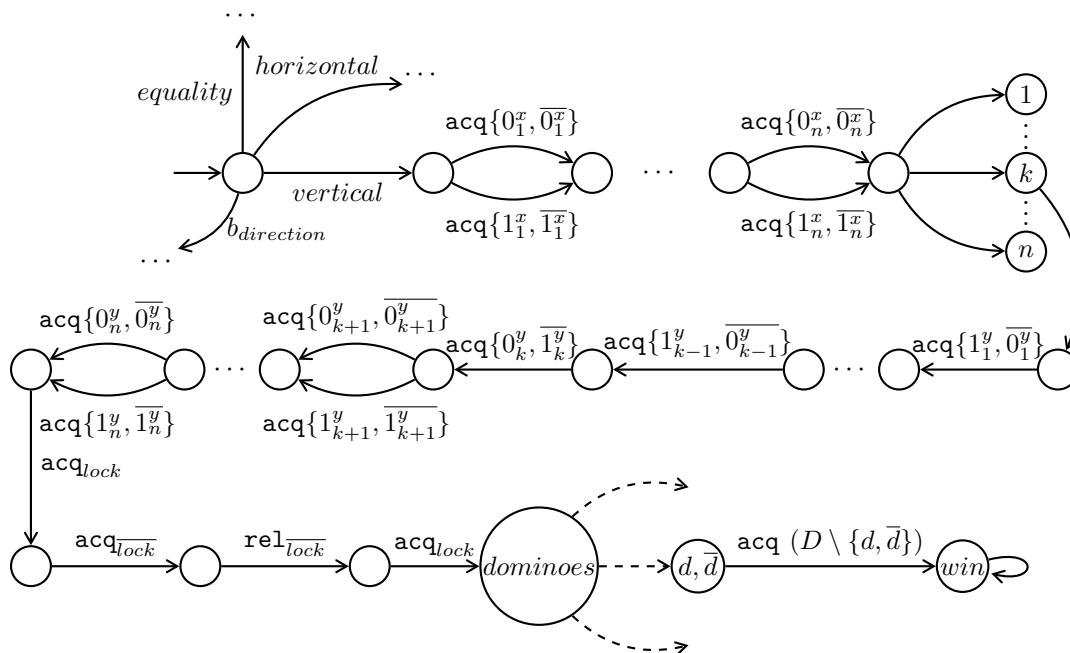


Figure 7 Transition system for process q , dashed arrows are controlled by the system, every state before *dominoes* has a self-loop that is not drawn and $\text{acq } S$ means that there is a sequence of forced transitions with the operations acq_t for each $t \in S$ (in some order). We only drew the subsystem used when the environment chooses *vertical*.

1455 E Undecidability for unrestricted lock-sharing systems

1456 E.1 Initial ownership of locks

1457 In a lock-sharing system all locks are assumed to be initially free. We consider now the
 1458 variant where some of the locks are initially owned by some processes.

1459 The input is a lock-sharing system $\mathcal{S} = ((\mathcal{A}_p)_{p \in \text{Proc}}, \Sigma^s, \Sigma^e, T)$ and an initial configuration
 1460 $C_{\text{init}} = (\text{init}_p, I_p)_{p \in \text{Proc}}$ with pairwise disjoint sets $I_p \subseteq T$. The question is whether there
 1461 exists a control strategy that guarantees that no run from C_{init} deadlocks.

1462 It turns out that this generalization of the deadlock-avoidance control problem is not
 1463 more difficult than our original problem, as stated in Lemma 35:

1464 ► **Lemma 35.** *There is a polynomial-time reduction from the control problem for lock-sharing
 1465 systems with initial configuration to the control problem where all locks are initially free. The
 1466 reduction adds $|\text{Proc}|$ new locks.*

1467 **Proof.** The system $\mathcal{S} = (\mathcal{A}_p)_{p \in \text{Proc}}, \Sigma^s, \Sigma^e, T)$ with initial ownership $(I_p)_{p \in \text{Proc}}$ is trans-

1468 formed into a new system \mathcal{S}_\emptyset with additional locks. The transformation introduces one extra
 1469 lock per process, denoted k_p and called *the key of p* . Each process uses in addition to T_p the
 1470 $|Proc|$ extra locks.

1471 The transition system \mathcal{A}_p of process p is extended with new states and transitions, which
 1472 define a specific finite run called the *init sequence*. The new states and transitions can occur
 1473 only during the init sequence. When a process p completes his init sequence in \mathcal{S}_\emptyset , he owns
 1474 precisely all locks in I_p , plus the key k_p , and has reached his initial state $init_p$ in \mathcal{A}_p . After
 1475 that, further actions and transitions played in \mathcal{S}_\emptyset are actions and transitions of \mathcal{S} , unchanged.
 1476 All the new actions are uncontrollable, thus there is no strategic decision to make for the
 1477 controller of a process p until his init sequence is completed.

1478 The init sequence.

1479 For process p , the init sequence consists of three steps.

- 1480 1. First, p takes one by one (in a fixed arbitrary order) all locks in I_p .
- 1481 2. Second, p takes and releases, one by one (in a fixed arbitrary order) all the keys of the
 1482 other processes $(k_q)_{q \neq p}$.
- 1483 3. Finally, p acquires his key k_p and keeps it forever.

1484 After acquiring k_p process p reaches the initial state $init_p$ in \mathcal{A}_p .

1485 In order to prevent the init sequence to create extra deadlocks, every state used in the
 1486 initialisation sequence is equipped with a local self-loop on the *nop* operation. This way, a
 1487 deadlock may only occur if all processes have finally completed their init sequences.

1488 Linking runs in \mathcal{S}_\emptyset and \mathcal{S} .

1489 When a process completes his init sequence, he has been until that point the sole owner of
 1490 its initial locks:

1491 \triangleright **Claim 55.** Let p be a process and u_\emptyset a run of \mathcal{S}_\emptyset such that the last action of u_\emptyset is \mathbf{acq}_{k_p}
 1492 by process p . Let $t \in I_p$, then p is the only process to acquire t in u_\emptyset .

1493 **Proof.** By contradiction, let $t \in I_p$ and $q \neq p$ and assume that u_\emptyset factorizes as $u_\emptyset =$
 1494 $u_0 \cdot (\mathbf{acq}_t, q) \cdot u_1 \cdot (\mathbf{acq}_{k_p}, p)$ (we abuse the notation and denote (\mathbf{acq}_t, q) and (\mathbf{acq}_{k_p}, p) the
 1495 transitions where q and p respectively acquire t and k_p). Process p must take and release k_q
 1496 before taking k_p , thus the transition $\delta = (\mathbf{acq}_{k_q}, p)$ occurs either in u_0 or in u_1 . However δ
 1497 cannot occur in u_0 : the init sequence of p requires that p owns t permanently in the interval
 1498 between the occurrence of δ and the occurrence of (\mathbf{acq}_{k_p}, p) , thus (\mathbf{acq}_t, q) cannot occur in
 1499 the meantime. Hence δ occurs in u_1 . But this leads to a contradiction: since t is not an
 1500 initial lock of q , process q is not allowed to acquire t during his init sequence, hence q has
 1501 completed his init sequence in u_0 . After u_0 , q owns permanently k_q , but then it is impossible
 1502 that $\delta = (\mathbf{acq}_{k_q}, p)$ occurs during u_1 . \blacktriangleleft

1503 There is a tight link between runs in \mathcal{S}_\emptyset and runs in \mathcal{S} .

1504 \triangleright **Claim 56.** Let u_\emptyset be a global run in \mathcal{S}_\emptyset in which all processes have completed their init
 1505 sequences. There exists a global run u in \mathcal{S} (with initial lock ownership $(I_p)_{p \in Proc}$) with the
 1506 same local runs as u_\emptyset , except that the init sequences are deleted.

1507 **Proof.** The proof is by induction on the number N of transitions in u_\emptyset which are not
 1508 transitions of the init sequence. In the base case $N = 0$, then u_\emptyset is an interleaving of the
 1509 init sequences of all processes and u is the empty run. Assume now $N > 0$. Let δ be the

1510 last transition played in u_\emptyset which is not part of an init sequence, and $Z \subseteq Proc$ the set of
 1511 processes that have not yet completed their init sequence when δ occurs. Then u_\emptyset factorizes
 1512 as

$$1513 \quad u_\emptyset = u'_\emptyset \cdot \delta \cdot u''_\emptyset$$

1514 where u''_\emptyset is an interleaving of infixes of the init sequences of processes in Z .

1515 Assume first that u''_\emptyset is empty. We apply the inductive hypothesis to u'_\emptyset , get a global run
 1516 ν and set $u = \nu \cdot \delta$. Then u has the same local runs as u_\emptyset , after deletion of init sequences.

1517 We now reduce the general case to the special case where u''_\emptyset is empty. Let q be the
 1518 process operating in δ and (a, op) the corresponding pair of action and operation on locks.
 1519 Since δ is not part of an init sequence, then $q \notin Z$ and op is not an operation on one of the
 1520 keys. Moreover, according to Claim 55, neither is op an operation on one of the initial locks
 1521 of processes in Z . Thus (a, op) can commute with all transitions in u''_\emptyset and become the last
 1522 transition of the global run, while leaving the local runs unchanged, and we are back to the
 1523 case where u''_\emptyset is empty. ◀

1524 We turn now to the proof of the theorem.

1525 ▷ **Claim 57.** The system wins in \mathcal{S}_\emptyset if and only if it wins in \mathcal{S} with initial lock ownership
 1526 $(I_p)_{p \in Proc}$.

1527 Since in \mathcal{S}_\emptyset there is no strategic decision to make during the init sequence, the strategies
 1528 in \mathcal{S}_\emptyset are in a natural one-to-one correspondence with strategies in \mathcal{S} . For a fixed strategy
 1529 we show that there is some deadlock in \mathcal{S} if and only if there is some deadlock in \mathcal{S}_\emptyset .

1530 If there is a deadlock in \mathcal{S} then there is also one in \mathcal{S}_\emptyset , by executing first all init sequences,
 1531 and then the deadlocking run of \mathcal{S} . The execution of all init sequences is in two steps: first
 1532 each process p acquires its initial locks I_p and acquires and releases the keys $k_q, q \neq p$ of
 1533 other processes. Second, each process p acquires (definitively) its key k_p .

1534 Suppose now that there is a deadlocking run u_\emptyset in \mathcal{S}_\emptyset . Observe first that all processes
 1535 $p \in Proc$ have completed their init sequences in u , because all states used in this sequence
 1536 have local *nop* self-loops. By Claim 56 there exists a global run u of \mathcal{S} which has the same
 1537 local runs as u_\emptyset (apart from the init sequences). Since u_\emptyset is deadlocking, so is u . ◀

1538 E.2 Undecidability

1539 In this section we show that the deadlock-avoidance control problem becomes undecidable if
 1540 we do not limit the maximal number of locks that processes can use.

1541 ► **Lemma 34.** *The control problem for lock-sharing systems with 3 processes, fixed initial*
 1542 *configuration and fixed number of locks per process is undecidable.*

1543 We reduce from the question whether a PCP instance has an infinite solution. Let
 1544 $(\alpha_i, \beta_i)_{i=1}^m$ be a PCP instance with $\alpha_i, \beta_i \in \{0, 1\}^*$. We construct below a system with three
 1545 processes P, \bar{P}, C , using locks from the set

$$1546 \quad \{c, s_0, s_1, p, \bar{s}_0, \bar{s}_1, \bar{p}\}.$$

1547 Process P will use locks from $\{c, s_0, s_1, p\}$, process \bar{P} from $\{c, \bar{s}_0, \bar{s}_1, \bar{p}\}$, and C all seven
 1548 locks.

1549 Processes P, \bar{P} are supposed to synchronize over a PCP solution with the controller
 1550 process C . That is, P and C synchronize over a sequence $\alpha_{i_1} \alpha_{i_2} \dots$, whereas \bar{P} and C

1551 synchronize over a sequence $\beta_{j_1}\beta_{j_2}\dots$. The environment tells C at the beginning whether
 1552 she should check index equality $i_1i_2\dots = j_1j_2\dots$ or word equality $\alpha_{i_1}\alpha_{i_2}\dots = \beta_{j_1}\beta_{j_2}\dots$.

1553 For the initial configuration we assume that P owns p , \bar{P} owns \bar{p} and C owns $c, s_0, s_1, \bar{s}_0, \bar{s}_1$.
 1554 We describe now the three processes P, \bar{P}, C . Define first for $b = 0, 1$:

$$1555 \quad u_P(b) = \text{acq}_{s_b} \text{rel}_p \text{acq}_c \text{rel}_{s_b} \text{acq}_p \text{rel}_c$$

$$1556 \quad u_{\bar{P}}(b) = \text{acq}_{\bar{s}_b} \text{rel}_{\bar{p}} \text{acq}_c \text{rel}_{\bar{s}_b} \text{acq}_{\bar{p}} \text{rel}_c$$

1557 The automaton of \mathcal{A}_P ($\mathcal{A}_{\bar{P}}$, resp.) allows all possible action sequences from $(u_P(0) +$
 1558 $u_P(1))^\omega$ ($(u_{\bar{P}}(0) + u_{\bar{P}}(1))^\omega$, resp.). If e.g. process P manages to execute a sequence
 1559 $u_P(b_1)u_P(b_2)\dots$ then this means that C, P synchronize over the sequence b_1, b_2, \dots .

1560 Process C 's behavior for checking word equality consists in repeating the following
 1561 procedure: she chooses a bit b through a controllable action, then tries to do $u_C(P, b)u_C(\bar{P}, b)$,
 1562 with:

$$1563 \quad u_C(P, b) = \text{rel}_{s_b} \text{acq}_p \text{rel}_c \text{acq}_{s_b} \text{rel}_p \text{acq}_c$$

$$1564 \quad u_C(\bar{P}, b) = \text{rel}_{\bar{s}_b} \text{acq}_{\bar{p}} \text{rel}_c \text{acq}_{\bar{s}_b} \text{rel}_{\bar{p}} \text{acq}_c$$

1565 For index equality C 's behavior is similar: she chooses an index i and then tries to do
 1566 $u_C(P, b_1)\dots u_C(P, b_k)u_C(\bar{P}, b'_1)\dots u_C(\bar{P}, b'_\ell)$, where $\alpha_i = b_1\dots b_k, \beta_i = b'_1\dots b'_\ell$.

1567 The next lemma is the key property showing that the system deadlock-avoiding strategy
 1568 if and only if the PCP instance has a solution.

1569 ► **Lemma 58.** *Assume that C owns $\{s_0, s_1, c\}$, P owns $\{p\}$, C wants to execute $u_C(P, b)$,
 1570 and P wants to execute $u_P(b')$. Then the system deadlocks if and only if $b \neq b'$. If $b = b'$
 1571 then C and P finish executing $u_C(P, b)$ and $u_P(b)$, respectively, and the lock ownership is the
 1572 same as before the execution.*

1573 **Proof.** If, say, $b = 0$ and $b' = 1$ then C releases s_0 but P wants to acquire s_1 , so that P
 1574 deadlocks. Since C wants to acquire p as second step, she deadlocks, too. Process \bar{P} will
 1575 deadlock as well, because he is waiting for either \bar{s}_0 or \bar{s}_1 .

1576 Suppose now that $b = b'$, say with $b = 0$. Then there is only one possible run alternating
 1577 between steps of $u_C(P, 0)$ and $u_P(0)$, until C finally acquires c . Then both C and P have
 1578 finished the execution of $u_C(P, 0)$ and $u_P(0)$, respectively. Moreover, C re-owns $\{c, s_0, s_1\}$
 1579 and P re-owns $\{p\}$. ◀