

# Synthesis from Component Libraries with Costs

Guy Avni and Orna Kupferman

School of Computer Science and Engineering, The Hebrew University, Israel

**Abstract.** *Synthesis* is the automated construction of a system from its specification. In real life, hardware and software systems are rarely constructed from scratch. Rather, a system is typically constructed from a library of components. Lustig and Vardi formalized this intuition and studied LTL synthesis from component libraries. In real life, designers seek optimal systems. In this paper we add optimality considerations to the setting. We distinguish between quality considerations (for example, size – the smaller a system is, the better it is), and pricing (for example, the payment to the company who manufactured the component). We study the problem of designing systems with minimal quality-cost and price. A key point is that while the quality cost is individual – the choices of a designer are independent of choices made by other designers that use the same library, pricing gives rise to a resource-allocation game – designers that use the same component share its price, with the share being proportional to the number of uses (a component can be used several times in a design). We study both closed and open settings, and in both we solve the problem of finding an optimal design. In a setting with multiple designers, we also study the game-theoretic problems of the induced resource-allocation game.

## 1 Introduction

Synthesis is the automated construction of a system from its specification. The classical approach to synthesis is to extract a system from a proof that the specification is satisfiable. In the late 1980s, researchers realized that the classical approach to synthesis is well suited to *closed* systems, but not to *open* (also called *reactive*) systems [1,24]. A reactive system interacts with its environment, and a correct system should have a *strategy* to satisfy the specification with respect to all environments. It turns out that the existence of such a strategy is stronger than satisfiability, and is termed *reliability*.

In spite of the rich theory developed for synthesis, in both the closed and open settings, little of this theory has been reduced to practice. This is in contrast with verification algorithms, which are extensively applied in practice. We distinguish between algorithmic and conceptual reasons for the little impact of synthesis in practice. The algorithmic reasons include the high complexity of the synthesis problem (PSPACE-complete in the closed setting [28] and 2EXPTIME-complete in the open setting [24], for specifications in LTL) as well as the intricacy of the algorithms in the open setting – the traditional approach involves determinization of automata on infinite words [27] and a solution of parity games [19].

We find the argument about the algorithmic challenge less compelling. First, experience with verification shows that even nonelementary algorithms can be practical,

since the worst-case complexity does not arise often. For example, while the model-checking problem for specifications in second-order logic has nonelementary complexity, the model-checking tool MONA [14] successfully verifies many specifications given in second-order logic. Furthermore, in some sense, synthesis is not harder than verification: the complexity of synthesis is given with respect to the specification only, whereas the complexity of verification is given with respect to the specification and the system, which is typically much larger than the specification. About the intercity of the algorithms, in the last decade we have seen quite many alternatives to the traditional approach – Safrless algorithms that avoid determinization and parity games, and reduce synthesis to problems that are simpler and are amenable to optimizations and symbolic implementations [17,21,22].

The arguments about the conceptual and methodological reasons are more compelling. We see here three main challenges, relevant in both the closed and open settings. First, unlike verification, where a specification can be decomposed into sub-specifications, each can be checked independently, in synthesis the starting point is one comprehensive specification. This inability to decompose or evolve the specification is related to the second challenge. In practice, we rarely construct systems from scratch or from one comprehensive specification. Rather, systems are constructed from existing components. This is true for both hardware systems, where we see IP cores or design libraries, and software systems, where web APIs and libraries of functions and objects are common. Third, while in verification we only automate the check of the system, automating its design is by far more risky and unpredictable – there are typically many ways to satisfy a satisfiable or realizable specification, and designers will be willing to give up manual design only if they can count on the automated synthesis tool to construct systems of comparable quality. Traditional synthesis algorithms do not attempt to address the quality issue.

In this paper we continue earlier efforts to cope with the above conceptual challenges. Our contribution extends both the setting and the results of earlier work. The realization that design of systems proceeds by composition of underlying components is not new to the verification community. For example, [18] proposed a framework for component-based modelling that uses an abstract layered model of components, and [12] initiated a series of works on interface theories for component-based design, possibly with a reuse of components in a library [13]. The need to consider components is more evident in the context of software, where, for example, recursion is possible, so components have to be equipped with mechanisms for call and return [4]. The setting and technical details, however, are different from these in the synthesis problem we consider here. The closer to our work here is [23], which studied LTL synthesis from reusable component libraries. Lustig and Vardi studied two notions of component composition. In the first notion, termed data-flow composition, components are cascaded so that the outputs of one component are fed to other components. In the second notion, termed control-flow composition, the composition is flat and control flows among the different components. The second notion, which turns out to be the decidable one [23], is particularly suitable for modelling web-service orchestration, where users are typically offered services and interact with different parties [3].

Let us turn now to the quality issue. Traditional formal methods are based on a Boolean satisfaction notion: a system satisfies, or not, a given specification. The richness of today’s systems, however, calls for specification formalisms that are *multi-valued*. The multi-valued setting arises directly in probabilistic and weighted systems and arises indirectly in applications where multi-valued satisfaction is used in order to model quantitative properties of the system like its size, security level, or quality. Reasoning about quantitative properties of systems is an active area of research in recent years, yielding quantitative specification formalisms and algorithms [11,16,10,2,9]. In quantitative reasoning, the Boolean satisfaction notion is refined and one can talk about the cost, or reward, of using a system, or, in our component-based setting, the cost of using a component from the library.

In order to capture a wide set of scenarios in practice, we associate with each component in the library two costs: a *quality cost* and a *construction cost*. The quality cost, as describes above, concerns the performance of the component and is paid each time the component is used. The construction cost is the cost of adding the component to the library. Thus, a design that uses a component pays its construction cost once. When several designs use the same component, they share its construction cost. This corresponds to real-life scenarios, where users pay, for example, for web-services, and indeed their price is influenced by the market demand.

We study synthesis from component libraries with costs in the closed and open settings. In both settings, the specification is given by means of a deterministic automaton  $\mathcal{S}$  on finite words (DFA).<sup>1</sup> In the closed setting, the specification is a regular language over some alphabet  $\Sigma$  and the library consists of box-DFAs (that is, DFAs with exit states) over  $\Sigma$ . In the open setting, the specification  $\mathcal{S}$  is over sets  $I$  and  $O$  of input and output signals, and the library consists of box- $I/O$ -transducers. The boxes are black, in the sense that a design that uses components from the library does not see  $\Sigma$  (or  $I \cup O$ ) nor it sees the behaviour inside the components. Rather, the mode of operation is as in the control-flow composition of [23]: the design gives control to one of the components in the library. It then sees only the exit state through which the component completes its computation and relinquishes control. Based on this information, the design decides which component gets control next, and so on.

In more technical details, the synthesis problem gets as input the specification  $\mathcal{S}$  as well as a library  $\mathcal{L}$  of components  $\mathcal{B}_1, \dots, \mathcal{B}_n$ . The goal is to return a correct design – a transducer  $\mathcal{D}$  that reads the exit states of the components and outputs the next component to gain control. In the closed setting, correctness means that the language over  $\Sigma$  that is generated by the composition defined by  $\mathcal{D}$  is equal to the language of  $\mathcal{S}$ . In the open setting, correctness means that the interaction of the composition defined by  $\mathcal{D}$  with all input sequences generates a computation over  $I \cup O$  that is in the language of  $\mathcal{S}$ .

We first study the problem without cost and reduce it to the solution of a two-player safety game  $\mathcal{G}_{\mathcal{L},\mathcal{S}}$ . In the closed setting, the game is of full information and the problem

---

<sup>1</sup> It is possible to extend our results to specifications in LTL. We prefer to work with deterministic automata, as this setting isolates the complexity and technical challenges of the design problem and avoids the domination of the doubly-exponential complexity of going from LTL to deterministic automata.

can be solved in polynomial time. In the open setting, the flexibility that the design have in responding to different input sequences introduces partial information to the game, and the problem is EXPTIME-complete. We note that in [23], where the open setting was studied and the specification is given by means of an LTL formula, the complexity is 2EXPTIME-complete, thus one could have expected our complexity to be only polynomial. We prove, however, hardness in EXPTIME, showing that it is not just the need to transfer the LTL formula to a deterministic formalism that leads to the high complexity.

We then turn to integrate cost to the story. As explained above, there are two types of costs associated with each component  $\mathcal{B}_i$  in  $\mathcal{L}$ . The first type, quality cost, can be studied for each design in isolation. We show that even there, the combinatorial setting is not simple. While for the closed setting an optimal design can be induced from a memoryless strategy of the designer in the game  $\mathcal{G}_{\mathcal{L},S}$ , making the problem of finding an optimal design NP-complete, seeking designs of optimal cost may require sophisticated compositions in the open setting. In particular, we show that optimal designs may be exponentially larger than other correct designs<sup>2</sup>, and that an optimal design may not be induced by a memoryless strategy in  $\mathcal{G}_{\mathcal{L},S}$ . We are still able to bound the size of an optimal transducer by the size of  $\mathcal{G}_{\mathcal{L},S}$ , and show that the optimal synthesis problem is NEXPTIME-complete.

The second type of cost, namely construction cost, depends not only on choices made by the designer, but also on choices made by designers of other specifications that use the library. Indeed, recall that the construction cost of a component is shared by designers that use this component, with the share being proportional to the number of uses (a component can be used several times in a design). Hence, the setting gives rise to a *resource-allocation game* [26,15]. Unlike traditional resource-allocation games, where players' strategies are sets of resources, here each strategy is a multiset – the components a designer needs. As has been the case in [7], the setting of multisets makes the game less stable. We show that the game is not guaranteed to have a *Nash Equilibrium* (NE), and that the problem of deciding whether an NE exists is  $\Sigma_2^P$ -complete. We then turn to the more algorithmic related problems and show that the problems of finding an optimal design given the choices of the other designers (a.k.a. the *best-response* problem, in algorithmic game theory) and of finding designs that minimize the total cost for all specifications (a.k.a. the *social optimum*) are both NP-complete.

## 2 Preliminaries

**Automata, transducers, and boxes** A *deterministic finite automaton* (DFA, for short) is a tuple  $\mathcal{A} = \langle \Sigma, Q, \delta, q_0, F \rangle$ , where  $\Sigma$  is an alphabet,  $Q$  is a set of states,  $\delta : Q \times \Sigma \rightarrow Q$  is a partial transition function,  $q_0 \in Q$  is an initial states, and  $F \subseteq Q$  is a set of accepting states. We extend  $\delta$  to words in an expected way, thus  $\delta^* : Q \times \Sigma^* \rightarrow Q$  is such that for  $q \in Q$ , we have  $\delta^*(q, \epsilon) = q$  and for  $w \in \Sigma^*$  and  $\sigma \in \Sigma$ , we have  $\delta^*(q, w \cdot \sigma) = \delta(\delta^*(q, w), \sigma)$ . When  $q = q_0$ , we sometimes omit it, thus  $\delta^*(w)$  is the state that  $\mathcal{A}$  reaches after reading  $w$ . We assume that all states are reachable from  $q_0$ , thus for every  $q \in Q$  there exists a word  $w \in \Sigma^*$  such that  $\delta^*(w) = q$ . We refer to the *size* of  $\mathcal{A}$ , denoted  $|\mathcal{A}|$ , as the number of its states.

<sup>2</sup> Recall that “optimal” here refers to the quality-cost function.

The *run* of  $\mathcal{A}$  on a word  $w = w_1, \dots, w_n \in \Sigma^*$  is the sequence of states  $r = r_0, r_1, \dots, r_n$  such that  $r_0 = q_0$  and for every  $0 \leq i \leq n - 1$  we have  $r_{i+1} = \delta(r_i, w_{i+1})$ . The run  $r$  is accepting iff  $r_n \in F$ . The *language* of  $\mathcal{A}$ , denoted  $L(\mathcal{A})$ , is the set of words  $w \in \Sigma^*$  such that the run of  $\mathcal{A}$  on  $w$  is accepting, or, equivalently,  $\delta^*(w) \in F$ . For  $q \in Q$ , we denote by  $L(\mathcal{A}^q)$  the language of the DFA that is the same as  $\mathcal{A}$  only with initial state  $q$ . Note that since  $\mathcal{A}$  is deterministic and  $\delta$  is partial, there is at most one run of  $\mathcal{A}$  on each word.

A *transducer* models an interaction between a system and its environment. It is similar to a DFA except that in addition to  $\Sigma$ , which is referred to as the input alphabet, denoted  $\Sigma_I$ , there is an output alphabet, denoted  $\Sigma_O$ , and rather than being classified to accepting or rejecting, each state is labeled by a letter from  $\Sigma_O$ <sup>3</sup>. Formally, a transducer is a tuple  $\mathcal{T} = \langle \Sigma_I, \Sigma_O, Q, q_0, \delta, \nu \rangle$ , where  $\Sigma_I$  is an input alphabet,  $\Sigma_O$  is an output alphabet,  $Q, q_0$ , and  $\delta : Q \times \Sigma_I \rightarrow Q$  are as in a DFA, and  $\nu : Q \rightarrow \Sigma_O$  is an output function. We require  $\mathcal{T}$  to be *receptive*. That is,  $\delta$  is complete, so for every input word  $w \in \Sigma_I^*$ , there is a run of  $\mathcal{T}$  on  $w$ . Consider an input word  $w = w_1, \dots, w_n \in \Sigma_I^*$ . Let  $r = r_0, \dots, r_n$  be the run of  $\mathcal{T}$  on  $w$ . The *computation* of  $\mathcal{T}$  in  $w$  is then  $\sigma_1, \dots, \sigma_n \in (\Sigma_I \times \Sigma_O)^*$ , where for  $1 \leq i \leq n$ , we have  $\sigma_i = \langle w_i, \nu(r_{i-1}) \rangle$ . We define the language of  $\mathcal{T}$ , denoted  $L(\mathcal{T})$ , as the set of all its computations. For a specification  $L \subseteq (\Sigma_I \times \Sigma_O)^*$ , we say that  $\mathcal{T}$  *realizes*  $L$  iff  $L(\mathcal{T}) \subseteq L$ . Thus, no matter what the input sequence is, the interaction of  $\mathcal{T}$  with the environment generates a computation that satisfies the specification. For two words  $u \in \Sigma_I^*$  and  $v \in \Sigma_O^*$  of length  $n$  we define the *product* of the two words, denoted  $u \oplus v$ , as  $w = w_1 \dots w_n \in (\Sigma_I \times \Sigma_O)^*$ , where, for  $1 \leq i \leq n$ , we have  $w_i = \langle u_i, v_i \rangle$ .

By adding *exit states* to DFAs and transducers, we can view them as components from which we can compose systems. Formally, we consider two types of components. Closed components are modeled by *box-DFAs* and open components are modeled by *box-transducers*. A box-DFA augments a DFA by a set of exit states. Thus, a box-DFA is a tuple  $\langle \Sigma, Q, \delta, q_0, F, E \rangle$ , where  $E \subseteq Q$  is a nonempty set of exit states. There are no outgoing transitions from an exit state. Also, the initial state cannot be an exit state and exit states are not accepting. Thus,  $q_0 \notin E$  and  $F \cap E = \emptyset$ . Box-transducers are defined similarly, and their exit states are not labeled, thus  $\nu : Q \setminus E \rightarrow \Sigma_O$ .

**Component libraries** A *component library* is a collection of boxes  $\mathcal{L} = \{\mathcal{B}_1, \dots, \mathcal{B}_n\}$ . We say that  $\mathcal{L}$  is a *closed library* if the boxes are box-DFAs, and is an *open library* if the boxes are box-transducers. Let  $[n] = \{1, \dots, n\}$ . In the first case, for  $i \in [n]$ , let  $\mathcal{B}_i = \langle \Sigma, C_i, \delta_i, c_i^0, F_i, E_i \rangle$ . In the second case,  $\mathcal{B}_i = \langle \Sigma_I, \Sigma_O, C_i, \delta_i, c_i^0, \nu_i, E_i \rangle$ . Note that all boxes in  $\mathcal{L}$  share the same alphabet (input and output alphabet, in the case of transducers). We assume that the states of the components are disjoint, thus for every  $i \neq j \in [n]$ , we have  $C_i \cap C_j = \emptyset$ . We use the following abbreviations  $\mathcal{C} = \bigcup_{i \in [n]} C_i$ ,  $\mathcal{C}_0 = \bigcup_{i \in [n]} \{c_i^0\}$ ,  $\mathcal{F} = \bigcup_{i \in [n]} F_i$ , and  $\mathcal{E} = \bigcup_{i \in [n]} E_i$ . We define the *size* of  $\mathcal{L}$  as  $|\mathcal{C}|$ .

We start by describing the intuition for composition of closed libraries. A *design* is a recipe to compose the components of a library  $\mathcal{L}$  (allowing multiple uses) into a DFA. A run of the design on a word starts in an initial state of one of the components in  $\mathcal{L}$ . We say that this component has the initial *control*. When a component is in control, the

<sup>3</sup> These transducers are sometimes referred to as *Moore machines*.

run uses its states, follows its transition function, and if the run ends, it is accepting iff it ends in one of the components' accepting states. A component relinquishes control when the run reaches one of its exit states. It is then the design's duty to assign control to the next component, which gains control through its initial state.

Formally, a design is a transducer  $\mathcal{D}$  with input alphabet  $\mathcal{E}$  and output alphabet  $[n]$ . We can think of  $\mathcal{D}$  as running beside the components. When a component reaches an exit state  $e$ , then  $\mathcal{D}$  reads the input letter  $e$ , proceeds to its next state, and outputs the index of the component to gain control next. Note that  $\mathcal{D}$  does not read the alphabet  $\Sigma$  and has no information about the states that the component visits. It only sees which exit state has been reached.

Consider a design  $\mathcal{D} = \langle \mathcal{E}, [n], D, \delta, d^0, \nu \rangle$  and a closed library  $\mathcal{L}$ . We formalize the behavior of  $\mathcal{D}$  by means of the *composition DFA*  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  that simulates the run of  $\mathcal{D}$  along with the runs of the box-DFAs. Formally,  $\mathcal{A}_{\mathcal{L}, \mathcal{D}} = \langle \Sigma, Q_{\mathcal{L}, \mathcal{D}}, \delta_{\mathcal{L}, \mathcal{D}}, q_{\mathcal{L}, \mathcal{D}}^0, F_{\mathcal{L}, \mathcal{D}} \rangle$  is defined as follows. The set of states  $Q_{\mathcal{L}, \mathcal{D}} \subseteq (\mathcal{C} \setminus \mathcal{E}) \times D$  consists of pairs of a *component state* from  $\mathcal{C}$  and an *design state* from  $S$ . The component states are consistent with  $\nu$ , thus  $Q_{\mathcal{L}, \mathcal{D}} = \bigcup_{i \in [n]} (C_i \setminus E_i) \times \{q : \nu(q) = i\}$ . In exit states, the composition immediately moves to the initial state of the next component, which is why the component states of  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  do not include  $\mathcal{E}$ . Consider a state  $\langle c, q \rangle \in Q_{\mathcal{L}, \mathcal{D}}$  and a letter  $\sigma \in \Sigma$ . Let  $i \in [n]$  be such that  $c \in C_i$ . When a run of  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  reaches the state  $\langle c, q \rangle$ , the component  $\mathcal{B}_i$  is in control. Recall that  $c$  is not an exit state. Let  $c' = \delta_i(c, \sigma)$ . If  $c' \notin E_i$ , then  $\mathcal{B}_i$  does not relinquish control after reading  $\sigma$  and  $\delta_{\mathcal{L}, \mathcal{D}}(\langle c, q \rangle, \sigma) = \langle c', q \rangle$ . If  $c' \in E_i$ , then  $\mathcal{B}_i$  relinquishes control through  $c'$ , and it is the design's task to choose the next component to gain control. Let  $q' = \delta(q, c')$  and let  $j = \nu(q')$ . Then,  $\mathcal{B}_j$  is the next component to gain control (possibly  $j = i$ ). Accordingly, we advance  $\mathcal{D}$  to  $q'$  and continue to the initial state of  $\mathcal{B}_j$ . Formally,  $\delta_{\mathcal{L}, \mathcal{D}}(\langle c, q \rangle, \sigma) = \langle c_j^0, q' \rangle$ . (Recall that  $c_j^0 \notin E_j$ , so the new state is in  $Q_{\mathcal{L}, \mathcal{D}}$ .) Note also that a visit in  $c'$  is skipped. The component that gains initial control is chosen according to  $\nu(d^0)$ . Thus,  $q_{\mathcal{L}, \mathcal{D}}^0 = \langle c_j^0, d^0 \rangle$ , where  $j = \nu(d^0)$ . Finally, the accepting states of  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  are these in which the component state is accepting, thus  $F_{\mathcal{L}, \mathcal{D}} = \mathcal{F} \times D$ .

The definition of a composition for an open library is similar. There, the composition is a transducer  $\mathcal{T}_{\mathcal{L}, \mathcal{D}} = \langle \Sigma_I, \Sigma_O, Q_{\mathcal{L}, \mathcal{D}}, \delta_{\mathcal{L}, \mathcal{D}}, q_{\mathcal{L}, \mathcal{D}}^0, \nu_{\mathcal{L}, \mathcal{D}} \rangle$ , where  $Q_{\mathcal{L}, \mathcal{D}}$ ,  $q_{\mathcal{L}, \mathcal{D}}^0$ , and  $\delta_{\mathcal{L}, \mathcal{D}}$  are as in the closed setting, except that  $\delta_{\mathcal{L}, \mathcal{D}}$  reads letters in  $\Sigma_I$ , and  $\nu_{\mathcal{L}, \mathcal{D}}(\langle c, q \rangle) = \nu_i(c)$ , for  $i \in [n]$  such that  $c \in C_i$ .

Consider a closed-library  $\mathcal{L}$ , a design  $\mathcal{D}$ , and the run  $r$  of  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  on  $w = w_0 \cdots w_l$ . We partition  $w$  according to positions in which control is transferred among components. Equivalently, positions in which  $r$  skips visits in exit states. Thus,  $w = y_0 \cdots y_k$  is such that for all  $0 \leq i < k$ , we have that  $y_i \in \Sigma^+$  and the composition  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  takes a transfer transition exactly when it reads the last letter of  $y_i$ . An exception is  $y_k$ , which may be empty (this happens when  $r$  ends upon entering the last component to gain control). We then say that  $w$  is *suffix-less*. The definitions in the open setting are similar.

### 3 The Design Problem

The *design problem* gets as input a component library  $\mathcal{L}$  and a specification that is given by means of a DFA  $\mathcal{S}$ . The problem is to decide whether there exists a correct

design for  $\mathcal{S}$  using the components in  $\mathcal{L}$ . In the closed setting, a design  $\mathcal{D}$  is correct if  $L(\mathcal{A}_{\mathcal{L},\mathcal{D}}) = L(\mathcal{S})$ . In the open setting,  $\mathcal{D}$  is correct if the transducer  $\mathcal{T}_{\mathcal{L},\mathcal{D}}$  realizes  $\mathcal{S}$ . Our solution to the design problem reduces it to the problem of finding the winner in a turn-based two-player game, defined below.

A *turn-based two-player game* is played on an arena  $\langle V, \Delta, V_0, \alpha \rangle$ , where  $V = V_1 \cup V_2$  is a set of vertices that are partitioned between Player 1 and Player 2,  $\Delta \subseteq V \times V$  is a set of directed edges,  $V_0 \subseteq V$  is a set of initial vertices, and  $\alpha$  is an objective for Player 1, specifying a subset of  $V^\omega$ . We consider here *safety games*, where  $\alpha \subseteq V$  is a set of vertices that are *safe* for Player 1. The game is played as follows. Initially, Player 1 places a token on one of the vertices in  $V_0$ . Assume the token is placed on a vertex  $v \in V$  at the beginning of a round. The player that owns  $v$  is the player that moves the token to the next vertex, where the legal vertices to continue to are  $\{v' \in V : \langle v, v' \rangle \in \Delta\}$ . The outcome of the game is a *play*  $\pi \in V^\omega$ . The play is winning for Player 1 if for every  $i \geq 1$ , we have  $\pi_i \in \alpha$ . Otherwise, Player 2 wins.

A *strategy* for Player  $i$ , for  $i \in \{1, 2\}$ , is a recipe that, given a prefix of a play, tells the player what his next move should be. Thus, it is a function  $f_i : V^* \cdot V_i \rightarrow V$  such that for every play  $\pi \cdot v \in V^*$  with  $v \in V_i$ , we have  $\langle v, f_i(\pi \cdot v) \rangle \in \Delta$ . Since Player 1 moves first, we require that  $f_1(\epsilon)$  is defined and is in  $V_0$ . For strategies  $f_1$  and  $f_2$  for players 1 and 2 respectively, the play  $out(f_1, f_2) \in V^\omega$  is the unique play that is the outcome the game when the players follow their strategies. A strategy  $f_i$  for Player  $i$  is *memoryless* if it depends only in the current vertex, thus it is a function  $f_i : V_i \rightarrow V$ .

A strategy is *winning* for a player if by using it he wins against every strategy of the other player. Formally, a strategy  $f_1$  is winning for Player 1 iff for every strategy  $f_2$  for Player 2, Player 1 wins the play  $out(f_1, f_2)$ . The definition for Player 2 is dual. It is well known that safety games are *determined*, namely, exactly one player has a winning strategy, and admits *memoryless* strategies, namely, Player  $i$  has a winning strategy iff he has a memoryless winning strategy. Deciding the winner of a safety game can done in linear time.

**Solving the design problem** We describe the intuition of our solution for the design problems. Given a library  $\mathcal{L}$  and a specification  $\mathcal{S}$  we construct a safety game  $\mathcal{G}_{\mathcal{L},\mathcal{S}}$  such that Player 1 wins  $\mathcal{G}_{\mathcal{L},\mathcal{S}}$  iff there is a correct design for  $\mathcal{S}$  using the components in  $\mathcal{L}$ . Intuitively, Player 1's goal is to construct a correct design, thus he chooses the components to gain control. Player 2 challenges the design that Player 1 chooses, thus he chooses a word (over  $\Sigma$  in the closed setting and over  $\Sigma_I \times \Sigma_O$  in the open setting) and wins if his word is a witness for the incorrectness of Player 1's design.

**Closed designs** The input to the closed-design problem is a closed-library  $\mathcal{L}$  and a DFA  $\mathcal{S}$  over the alphabet  $\Sigma$ . The goal is to find a correct design  $\mathcal{D}$ . Recall that  $\mathcal{D}$  is correct if the DFA  $\mathcal{A}_{\mathcal{L},\mathcal{D}}$  that is constructed from  $\mathcal{L}$  using  $\mathcal{D}$  satisfies  $L(\mathcal{A}_{\mathcal{L},\mathcal{D}}) = L(\mathcal{S})$ . We assume that  $\mathcal{S}$  is the minimal DFA for the language  $L(\mathcal{S})$ .

**Theorem 1.** *The closed-design problem can be solved in polynomial time.*

**Proof:** Given a closed-library  $\mathcal{L}$  and a DFA  $\mathcal{S} = \langle \Sigma, S, \delta_S, s^0, F_S \rangle$ , we describe a safety game  $\mathcal{G}_{\mathcal{L},\mathcal{S}}$  such that Player 1 wins  $\mathcal{G}_{\mathcal{L},\mathcal{S}}$  iff there is a design of  $\mathcal{S}$  using com-

ponents from  $\mathcal{L}$ . Recall that  $\mathcal{L}$  consists of box-DFAs  $\mathcal{B}_i = \langle \Sigma, C_i, \delta_i, c_i^0, F_i, E_i \rangle$ , for  $i \in [n]$ , and that we use  $\mathcal{C}$ ,  $\mathcal{C}_0$ ,  $\mathcal{E}$ , and  $\mathcal{F}$  to denote the union of all states, initial states, exit states, and accepting states in all the components of  $\mathcal{L}$ . The number of vertices in  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$  is  $|(\mathcal{C}_0 \cup \mathcal{E}) \times \mathcal{S}|$  and it can be constructed in polynomial time. Since solving safety games can be done in linear time, the theorem follows.

We define  $\mathcal{G}_{\mathcal{L}, \mathcal{S}} = \langle V, E, V_0, \alpha \rangle$ . First,  $V = (\mathcal{C}_0 \cup \mathcal{E}) \times \mathcal{S}$  and  $V_0 = \mathcal{C}_0 \times \{s^0\}$ . Recall that Player 1 moves when it is time to decide the next (or first) component to gain control. Accordingly,  $V_1 = \mathcal{E} \times \mathcal{S}$ . Also, Player 2 challenges the design suggested by Player 1 and chooses the word that is processed in a component that gains control, so  $V_2 = \mathcal{C}_0 \times \mathcal{S}$ .

Consider a vertex  $\langle e, s \rangle \in V_1$ . Player 1 selects the next component to gain control. This component gains control through its initial state. Accordingly,  $E$  contains edges  $\langle \langle e, s \rangle, \langle c_i^0, s \rangle \rangle$ , for every  $i \in [n]$ . Note that since no letter is read when control is passed, we do not advance the state in  $\mathcal{S}$ . Consider a vertex  $v = \langle c_i^0, s \rangle \in V_2$ . Player 2 selects the word that is read in the component  $\mathcal{B}_i$ , or equivalently, he selects the exit state from which  $\mathcal{B}_i$  relinquishes control. Thus,  $E$  contains an edge  $\langle \langle c_i^0, s \rangle, \langle e, s' \rangle \rangle$  iff there exists a word  $u \in \Sigma^*$  such that  $\delta_i^*(u) = e$  and  $\delta_{\mathcal{S}}^*(s, u) = s'$ .

We now turn to define the winning condition. All the vertices in  $V_1$  are in  $\alpha$ . A vertex  $v \in V_2$  is not in  $\alpha$  if it is possible to extend the word traversed for reaching  $v$  to a witness for the incorrectness of  $\mathcal{D}$ . Accordingly, a vertex  $\langle c_i^0, s \rangle$  is not in  $\alpha$  if one of the following holds. First (“the suffix witness”), there is a finite word that is read inside the current component and witnesses the incorrectness. Formally, there is  $u \in \Sigma^*$  such that  $\delta_i^*(u) \in F_i$  and  $\delta_{\mathcal{S}}^*(s, u) \notin F_{\mathcal{S}}$ , or  $\delta_i^*(u) \in C_i \setminus (F_i \cup E_i)$  and  $\delta_{\mathcal{S}}^*(s, u) \in F_{\mathcal{S}}$ . Second (“the infix witness”), there are two words that reach the same exit state of the current component yet the behavior of  $\mathcal{S}$  along them is different. Formally, there exist words  $u, u' \in \Sigma^*$  such that  $\delta_i^*(u) = \delta_i^*(u') \in E_i$  and  $\delta_{\mathcal{S}}^*(s, u) \neq \delta_{\mathcal{S}}^*(s, u')$ . Intuitively, the minimality of  $\mathcal{S}$  enables us to extend either  $u$  or  $u'$  to an incorrectness witness. Given  $\mathcal{L}$  and  $\mathcal{S}$ , the game  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$  can be constructed in polynomial time.

We claim that there is a correct design  $\mathcal{D}$  iff Player 1 wins  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$ . Assume first that there is a correct design  $\mathcal{D} = \langle \mathcal{E}, [n], D, \delta, d^0, \nu \rangle$ , thus  $L(\mathcal{A}_{\mathcal{L}, \mathcal{D}}) = L(\mathcal{S})$ . We construct a winning strategy  $f_{\mathcal{D}}$  for Player 1. The strategy  $f_{\mathcal{D}}$  proceeds like  $\mathcal{D}$ . First,  $f_{\mathcal{D}}(\epsilon) = \langle c_i^0, s^0 \rangle$ , with  $i = \nu(d^0)$ . Then, for a finite play  $\pi$ , let  $\langle e_0, s_0 \rangle, \langle e_1, s_1 \rangle \dots \langle e_m, s_m \rangle$  be its projection on  $V_1$ . Thus,  $e_0, \dots, e_m \in \mathcal{E}$  and  $s_0, \dots, s_m \in \mathcal{S}$ . We define  $f_{\mathcal{D}}(\pi) = \langle c_i^0, s_m \rangle$ , for  $i = \nu(\delta_{\mathcal{D}}^*(e_0, \dots, e_m))$ .

We claim that  $f_{\mathcal{D}}$  is a winning strategy. Assume towards contradiction that there is a Player 2 strategy  $f_2$  that wins against  $f_{\mathcal{D}}$ . Let  $\pi = \text{out}(f_{\mathcal{D}}, f_2)$ , let  $\langle c_i^0, s \rangle$  be its first vertex that is not in  $\alpha$ , and let  $w$  be the word in  $\Sigma^*$  that Player 2 follows along the prefix of  $\pi$  that reaches  $\langle c_i^0, s \rangle$ . Finally, let  $d \in D$  be such that  $\langle c_i^0, d \rangle$  is the state that  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  reaches when it reads  $w$ . Since we define  $f_{\mathcal{D}}$  to agree with  $\mathcal{D}$ , then the component state of this state in  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  is indeed  $c_i^0$ .

We distinguish between two cases. First, if  $\langle c_i^0, s \rangle$  exits  $\alpha$  because of a suffix witness, let  $u \in \Sigma^*$  be such that  $\delta_i^*(u) \in F_i$  and  $\delta_{\mathcal{S}}^*(s, u) \notin F_{\mathcal{S}}$ . (The case where  $\delta_i^*(u) \notin (F_i \cup E_i)$  and  $\delta_{\mathcal{S}}^*(s, u) \in F_{\mathcal{S}}$  is similar). Let  $c = \delta_i^*(u)$ . The run of  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  on  $w \cdot u$  ends in the state  $\langle \delta_i^*(u), s \rangle$ . Since  $\delta_i^*(u) \in F_i$ , we have  $w \cdot u \in L(\mathcal{A}_{\mathcal{L}, \mathcal{D}})$ . By the



definition of  $E$ , we have that  $\delta_S^*(w) = s$ . Since  $\delta_S^*(s, u) \notin F_S$ , we have  $w \cdot u \notin L(S)$ . Thus,  $L(\mathcal{A}_{\mathcal{L}, \mathcal{D}}) \neq L(S)$ , and we reach a contradiction to the correctness of  $\mathcal{D}$ .

In the second case, of an infix witness, there exist words  $u, u' \in \Sigma^*$  such that  $\delta_i^*(u) = \delta_i^*(u') \in E_i$  and  $\delta_S^*(s, u) = p \neq p' = \delta_S^*(s, u')$ . Since  $\mathcal{S}$  is a minimal DFA for  $L(S)$ , we have  $L(S^p) \neq L(S^{p'})$ . Thus, wlog, there is a word  $z \in L(S^p) \setminus L(S^{p'})$ . Recall that  $\delta_{\mathcal{L}, \mathcal{D}}^*(w) = \langle c_i^0, d \rangle$ . Let  $d' = \delta_{\mathcal{D}}(d, e)$  and  $j = \nu(d')$ . Since  $\delta_{\mathcal{L}, \mathcal{D}}^*(w \cdot u) = \delta_{\mathcal{L}, \mathcal{D}}^*(w \cdot u') = \langle c_j^0, d' \rangle$ , we have  $\delta_{\mathcal{L}, \mathcal{D}}^*(w \cdot u \cdot z) = \delta_{\mathcal{L}, \mathcal{D}}^*(w \cdot u' \cdot z)$ . Thus,  $w \cdot u \cdot z \in L(\mathcal{A}_{\mathcal{L}, \mathcal{D}})$  iff  $w \cdot u' \cdot z \in L(\mathcal{A}_{\mathcal{L}, \mathcal{D}})$ . However,  $w \cdot u \cdot z \in L(S)$  and  $w \cdot u' \cdot z \notin L(S)$ . Thus, we reach a contradiction to the correctness of  $\mathcal{D}$ , and we are done.

Assume now that Player 1 wins the game  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$ . Let  $f$  be a memoryless winning strategy for Player 1. We construct a correct design  $\mathcal{D}_f$  from  $f$ . Note that all the successors of a vertex in  $V_1$  are in  $V_2$ . Thus,  $f : V_1 \rightarrow V_2$ . We define  $\mathcal{D}_f = \langle \mathcal{E}, [n], D, \delta, s^0, \nu \rangle$  as follows. First,  $D = V_2 = \mathcal{C}_0 \times S$ . Consider a state  $v = \langle c_i^0, s \rangle \in V_2 \cap \alpha$ . Recall that  $c_i^0$  is the initial state of the component  $\mathcal{B}_i$ . Since  $v \in \alpha$ , the lack of an infix witness implies that for every exit state  $e \in E_i$  there is exactly one state  $s' \in S$  such that  $\langle \langle c_i^0, s \rangle, \langle e, s' \rangle \rangle \in E$ . We define  $\delta(v, e) = f(\langle e, s' \rangle)$ . Note that if  $v \notin \alpha$  or  $e \notin E_i$ , then we can define  $\delta(v, e)$  arbitrarily. The labeling function  $\nu$  is defined as expected, with  $\nu(\langle c_i^0, s \rangle) = i$ .

We prove that  $\mathcal{D}_f$  is a correct design. Assume towards contradiction that there is a word  $w \in L(\mathcal{A}_{\mathcal{L}, \mathcal{D}_f}) \setminus L(S)$ . The case where  $w \in L(S) \setminus L(\mathcal{A}_{\mathcal{L}, \mathcal{D}_f})$  is similar. Consider the run  $r$  of  $\mathcal{A}_{\mathcal{L}, \mathcal{D}_f}$  on  $w$ . Let  $\mathcal{B}_{i_1}, \dots, \mathcal{B}_{i_m} \in \mathcal{L}^*$  be sequence of components that  $r$  traverses and  $e_{i_1}, \dots, e_{i_{m-1}} \in \mathcal{E}^*$  be the corresponding exit states. Let  $y_1, \dots, y_m$  be the partition of  $w$  according to  $\mathcal{D}$ . Thus, for  $1 \leq j < m$ , we have  $y_j \in \Sigma^+$  and  $\delta_{i_j}^*(y_j) = e_{i_j} \in E_{i_j}$ , and  $y_m \in \Sigma^*$ . Since  $w \in L(\mathcal{A}_{\mathcal{L}, \mathcal{D}_f})$ , we have  $\delta_{i_m}^*(w_m) \in F_{i_m}$ . Note that the word  $y_1 \dots y_{m-1} \in \Sigma^*$  is suffix-less, thus  $\delta_{\mathcal{L}, \mathcal{D}_f}^*(y_1 \dots y_{m-1}) = \langle c_{i_m}^0, d \rangle$  for some  $d \in D$  with  $\nu(d) = i_m$ . Since we defined  $\mathcal{D}_f$  to agree with  $f$  on the components that gain control, the finite play  $\pi$  that is the outcome of the game when Player 1 plays  $f$  and Player 2 chooses the exit states  $e_{i_1}, \dots, e_{i_{m-1}}$  reaches the Player 2 vertex  $v = \langle c_{i_m}^0, s \rangle \in V_2$ , for  $s \in S$  such that  $\delta_S^*(y_1 \dots y_{m-1}) = s$ . We claim that  $v \notin \alpha$ . Indeed,  $\delta_i^*(y_m) \in F_i$  and since  $w \notin L(S)$ , we have  $y_m \notin L(S^s)$ . Thus,  $\pi$  is a winning play for Player 2, contradicting our assumption that  $f$  is a winning strategy, and we are done.  $\blacksquare$

**Open designs** We continue to study the open setting. Recall that there, the input is a DFA  $\mathcal{S}$  over the alphabet  $\Sigma_I \times \Sigma_O$  and an open library  $\mathcal{L}$ . The goal is to find a correct design  $\mathcal{D}$  or return that no such design exists, where  $\mathcal{D}$  is correct if the composition transducer  $\mathcal{T}_{\mathcal{L}, \mathcal{D}}$  realizes  $L(\mathcal{S})$ .

Lustig and Vardi [23] studied the design problem in a setting in which the specification is given by means of an LTL formula. They showed that the problem is 2EXPTIME-complete. Given an LTL formula one can construct a deterministic parity automaton that recognizes the language of words that satisfy the formula. The size of the automaton is doubly-exponential in the size of the formula. Thus, one might guess that the design problem in a setting in which the specification is given by means of a DFA would be solvable in polynomial time. We show that this is not the case and that the problem is EXPTIME-complete. As in [23], our upper bound is based on the ability to “summa-

size” the activity inside the components. Starting with an LTL formula, the solution in [23] has to combine the complexity involved in the translation of the LTL formula into an automaton with the complexity of finding a design, which is done by going through-out a universal word automaton that is expanded to a tree automaton. Starting with a deterministic automaton, our solution directly uses games. The interesting contribution, however, is the lower bound, showing that problem is EXPTIME-hard even when the specification is given by means of a deterministic automaton. We start with the upper bound.

**Theorem 2.** *The open-design problem is in EXPTIME.*

**Proof:** Given an open-library  $\mathcal{L}$  and a DFA  $\mathcal{S} = \langle \Sigma_I \times \Sigma_O, S, \delta_S, s^0, F_S \rangle$ , we describe a safety game  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$  such that Player 1 wins  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$  iff there is a design for  $\mathcal{S}$  using components from  $\mathcal{L}$ . The number of vertices in  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$  is exponential in  $S$  and  $\mathcal{C}$ . Since solving safety games can be done in linear time, membership in EXPTIME follows.

We define  $\mathcal{G}_{\mathcal{L}, \mathcal{S}} = \langle V, E, V_0, \alpha \rangle$  as follows.<sup>4</sup> Recall that  $\mathcal{C}$ ,  $\mathcal{C}_0$ ,  $\mathcal{E}$ , and  $\mathcal{F}$  are the union of all states, initial states, exit states, and accepting states in all the components of  $\mathcal{L}$ . We define  $V = (\mathcal{C}_0 \cup \mathcal{E}) \times 2^S$  with  $V_1 = \mathcal{E} \times 2^S$  and  $V_2 = \mathcal{C}_0 \times 2^S$ . Also,  $V_0 = \mathcal{C}_0 \times \{\{s^0\}\}$ . As in the closed-setting, Player 1 selects the components that gain control, thus for a vertex  $\langle e, T \rangle \in V_1$  we have  $\langle \langle e, T \rangle, \langle c_i^0, T \rangle \rangle \in E$ , for every  $c_i^0 \in \mathcal{C}_0$ . Player 2 selects the word that is processed in the component, or equivalently, the exit state from which it relinquishes control, thus for a vertex  $v = \langle c_i^0, T \rangle \in V_2$  we have  $\langle \langle c_i^0, T \rangle, \langle e, T' \rangle \rangle \in E$  iff for every  $s' \in T'$  there is a state  $s \in T$  and a word  $u \in \Sigma_I^*$  such that  $\delta_i^*(u) = e$  and, assuming  $w \in (\Sigma_I \times \Sigma_O)^*$  is the computation of  $\mathcal{B}_i$  that corresponds to  $u$ , we have  $\delta_S^*(s, w) = s'$ . Note that for  $c_i^0 \in \mathcal{C}_0$ ,  $T \in 2^S$ , and  $e \in \mathcal{E}$ , there is at most one, nonempty, subset  $T' \in 2^S$  such that  $\langle \langle c_i^0, T \rangle, \langle e, T' \rangle \rangle \in E$ . The set of vertices that are losing for Player 1 consists of states  $\langle c_i^0, T \rangle \in V_2$  from which Player 2 can generate a suffix-witness to the incorrectness of the design. Formally,  $\langle c_i^0, T \rangle \in \alpha$  iff there exists  $u \in L(\mathcal{B}_i)$  and  $s \in T$  such that  $u \notin L(\mathcal{S}^s)$ .

We claim that there is a correct design  $\mathcal{D}$  iff Player 1 wins  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$ . For the first direction, consider a correct design  $\mathcal{D}$ , thus  $L(\mathcal{T}_{\mathcal{L}, \mathcal{D}}) \subseteq L(\mathcal{S})$ . We construct a winning Player 1 strategy  $f_{\mathcal{D}}$ . Recall that a design reads exit states and outputs components. Further recall that assuming the game does not end, a Player 2 move is a choice of an exit state. We define  $f_{\mathcal{D}}$  so that it responds to Player 2’s choice the same way  $\mathcal{D}$  responds. We define  $f_{\mathcal{D}}(\epsilon) = \langle c_i^0, \{\{s^0\}\} \rangle$  for  $i = \nu(d^0)$ , thus  $f_{\mathcal{D}}$  and  $\mathcal{D}$  assign initial control to the same component. Consider a finite play  $\pi$  and let  $\langle e_1, T_1 \rangle, \dots, \langle e_m, T_m \rangle$  be the projection of  $\pi$  on  $V_2$ , thus  $e_1, \dots, e_m \in \mathcal{E}$  and  $T_1, \dots, T_m \in 2^S$ . We define  $f_{\mathcal{D}}(\pi) = \langle c_i^0, T_m \rangle$  where  $\nu(\delta_{\mathcal{D}}^*(e_1, \dots, e_m)) = i$ .

We claim that  $f_{\mathcal{D}}$  is a winning strategy. Assume towards contradiction that there is a Player 2 strategy  $f_w$  that wins against  $f_{\mathcal{D}}$ . Let  $\pi = \langle c_{i_1}^0, \{\{s^0\}\} \rangle, \langle e_{i_1}, T_1 \rangle, \langle c_{i_2}^0, T_1 \rangle, \langle e_{i_2}, T_2 \rangle, \dots, \langle e_{i_{m-1}}, T_m \rangle, \langle c_{i_m}^0, T_m \rangle$  be the finite losing prefix of  $out(f_{\mathcal{D}}, f_w)$ , thus  $\langle c_{i_m}^0, T_m \rangle \notin \alpha$ . Since  $\langle c_{i_m}^0, T_m \rangle \notin \alpha$  there is a state  $s_m \in T_m$  and a word  $w_m \in L(\mathcal{B}_i)$

<sup>4</sup> A different way to construct  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$  would be to go through a *partial-information* game (see Theorem 3) with vertices in  $\mathcal{C} \times S$ , where Player 1 cannot distinguish between vertices  $\langle e, d \rangle$  and  $\langle e, d' \rangle$  for  $e \in \mathcal{E}$  and  $d, d' \in S$ . The game  $\mathcal{S}$  is the corresponding game. We describe it directly, which also shows that the exponential dependency is only in  $S$ .

such that  $w_m \notin L(\mathcal{S}^{s_m})$ . It is not hard to see that there are words  $w_1, \dots, w_{m-1} \in (\Sigma_I \times \Sigma_O)^*$  and states  $s_1, \dots, s_{m-1} \in S$  such that for  $1 \leq j \leq m-1$  we have  $w_j \in L(\mathcal{B}_{i_j})$ ,  $\delta_{i_j}^*(c_{i_j}^0, w_j) = e_{i_j}$ ,  $\delta_S^*(w_1) = s_1$ , and  $\delta_S^*(s_j, w_j) = s_{j+1}$ . It is not hard to see that since we defined  $f_{\mathcal{D}}$  to agree with  $\mathcal{D}$ , the components that gain control in the run of  $\mathcal{T}_{\mathcal{L}, \mathcal{D}}$  on  $w = w_1 \cdot \dots \cdot w_m$  are  $\mathcal{B}_{i_1}, \dots, \mathcal{B}_{i_m}$ , thus it is possible to prove by induction that  $w \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}})$ . Moreover, the run of  $\mathcal{S}$  on  $w$  is not accepting, thus  $w \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}}) \setminus L(\mathcal{S})$ , and we reach a contradiction to the correctness of  $\mathcal{D}$ .

We continue to the second direction. Assume Player 1 wins the game  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$ . Thus, he has a memoryless winning strategy  $f_{\mathcal{D}}$  from which we construct a design  $\mathcal{D}$ . Intuitively, in  $\mathcal{D}$ , we skip exit states and proceed according to  $f_{\mathcal{D}}$ . Thus, the states of  $\mathcal{D}$  are  $V_2 = \mathcal{C}_0 \times 2^S$ . Consider a state  $\langle c_i^0, T \rangle \in V_2$ , where recall that  $c_i^0$  is the initial state of the component  $\mathcal{B}_i \in \mathcal{L}$ , and an exit state  $e \in E_i$  of  $\mathcal{B}_i$ . Recall that there is a unique subset  $T' \in 2^S$  such that  $\langle \langle c_i^0, T \rangle, \langle e, T' \rangle \rangle \in E$ . We define  $\delta_{\mathcal{D}}(\langle c_i^0, T \rangle, e) = f_{\mathcal{D}}(\langle e, T' \rangle)$ .

We claim that  $\mathcal{D}$  is a correct design. Assume towards contradiction that there is a word  $w \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}}) \setminus L(\mathcal{S})$ . Consider the run  $r$  of  $\mathcal{T}_{\mathcal{L}, \mathcal{D}}$  on  $w$ . Let  $\mathcal{B}_{i_1}, \dots, \mathcal{B}_{i_m} \in \mathcal{L}^*$  be sequence of components that  $r$  traverses. Let  $w_1, \dots, w_m$  be the partition of  $w$  according to  $\mathcal{D}$ . That is, for  $1 \leq j \leq m$ , the subword  $w_j$  is induced while  $r$  is in component  $\mathcal{B}_{i_j}$  and  $\delta_{i_j}^*(w_j) \in E_{i_j}$ . Let  $\pi_w = e_{i_1}, \dots, e_{i_{m-1}} \in \mathcal{E}^*$  be the exit states that  $r$  visits. Note that since  $\mathcal{B}_{i_m}$  gains control last the word  $w_1 \cdot \dots \cdot w_{m-1}$  is suffix-less, thus  $\delta_{\mathcal{L}, \mathcal{D}}^*(w_1 \cdot \dots \cdot w_{m-1}) = \langle c_{i_m}^0, d \rangle$  for some  $d \in D$  having  $\nu(d) = i_m$ . Moreover,  $w_m \in L(\mathcal{B}_{i_m})$ . Since we defined  $\mathcal{D}$  to agree with  $f_{\mathcal{D}}$  on the components that gain control, the finite play  $\pi$  that is the outcome of the game when Player 1 plays  $f_{\mathcal{D}}$  and Player 2 chooses the exit states  $e_{i_1}, \dots, e_{i_{m-1}}$  reaches the Player 2 vertex  $v = \langle c_{i_m}^0, T \rangle \in V_2$ , for some  $T \in 2^S$ . We claim that  $v \notin \alpha$ . Indeed, the definition of  $E$  implies that there is a vertex  $s \in T$  such that  $\delta_S^*(w_1 \cdot \dots \cdot w_{m-1}) = s$ . Since  $w \notin L(\mathcal{S})$ , we have  $w_m \notin L(\mathcal{S}^s)$  and  $w_m \in L(\mathcal{B}_{i_m})$ . Thus,  $\pi$  is a winning play for Player 2, contradicting our assumption that  $f_{\mathcal{D}}$  is a winning strategy, and we are done.  $\blacksquare$

We continue to study the lower bound.

**Theorem 3.** *The open-design problem is EXPTIME-hard.*

**Proof:** We describe a reduction from the problem of deciding whether Player 1 has a winning strategy in a *partial-information safety game*, known to be EXPTIME-complete [8].

Partial-information games (PI games, for short) are a variant of the *full-information* games (FI games, for short) defined above in which Player 1 has imperfect information [25]. The vertices, which we refer to as *locations*, denoted  $L$ , are partitioned into *observations*, denoted  $\mathcal{O}$ . Player 1 is unaware of the location on which the token is placed and is only aware of the observation it is in. Accordingly, In his turn, Player 1 cannot select the next location to move the token to. Instead, the edges in the game are labeled with actions, denoted  $\Gamma$ . In each round, Player 1 selects an action and Player 2 resolves nondeterminism and chooses the next location the token moves to. Initially, the token is placed on  $l^0 \in L$ . The set of labeled edges in the game is  $A \subseteq L \times \Gamma \times L$ , and the safety objective  $\alpha$  is given with respect to the observations, thus  $\alpha \subseteq \mathcal{O}$ .

Formally, a PI game is played on an arena  $\langle \Gamma, L, \mathcal{O}, l_0, A, \alpha \rangle$ , where  $\Gamma$  is a set of actions,  $L$  is a set of locations,  $\mathcal{O} \subseteq 2^L$  is a set of observations that form a partition of

$L$  (that is,  $\mathcal{O} = \{L_1, \dots, L_k\}$  where for all  $i \neq j \in [k]$ , we have that  $L_i \cap L_j = \emptyset$ , and  $\bigcup_{i \in [k]} L_i = L$ ),  $l_0 \in L$  is an initial location,  $\Lambda \subseteq L \times \Gamma \times L$  are labeled edges, and  $\alpha$  is a winning condition defined with respect to  $\mathcal{O}$ . In particular, in safety games  $\alpha \subseteq \mathcal{O}$  are safe observations for Player 1. We require that for every location  $l \in L$  and action  $a \in \Gamma$ , there is a location  $l' \in L$  such that  $\langle l, a, l' \rangle \in \Lambda$ .

The game proceeds similarly to FI games. At the beginning of the game, a token is placed on the initial location  $l_0$ . Assume the token is on a location  $l \in L$ . Player 1 moves first and selects an action  $a \in \Gamma$ . Player 2 resolves non-determinism and selects a location  $l' \in L$  such that  $\langle l, a, l' \rangle \in \Lambda$ . Since Player 1 only observes the member in  $\mathcal{O}$  in which  $l$  is, a strategy for Player 1 is a function  $f_1 : (\mathcal{O} \cdot \Gamma)^* \cdot \mathcal{O} \rightarrow \Gamma$ . Since Player 2 has complete information, a strategy for him is a function  $f_2 : L^+ \cdot \Gamma \rightarrow L$  such that for a play  $\pi = l_0, \dots, l_m$  and an action  $a \in \Gamma$ , we have  $\langle l_m, a, f_2(\pi, a) \rangle \in \Lambda$ . The definition of winning strategies are as in the full-information setting.

Consider a PI safety game  $\mathcal{G} = \langle \Gamma, L, \mathcal{O}, l^0, \Lambda, \alpha \rangle$ . We construct a library  $\mathcal{L}$  and a DFA  $\mathcal{S}$  such that there is a correct design for  $\mathcal{S}$  using the components of  $\mathcal{L}$  iff Player 1 wins  $\mathcal{G}$ . Recall that a design reads an exit state and outputs the component that gets control next. Also, a Player 1 strategy in  $\mathcal{G}$  reads observations and outputs actions. Accordingly, the library  $\mathcal{L}$  consists of box-transducers  $\mathcal{B}_a$ , one for every action  $a \in \Gamma$ . The exit states of the components correspond to the observations in  $\mathcal{O}$ . That is, when a component exits through an observation  $L_i \in \mathcal{O}$ , the design decides which component  $\mathcal{B}_a \in \mathcal{L}$  gains control, which corresponds to a Player 1 strategy that chooses the action  $a \in \Gamma$  from the observation  $L_i$ .

Next, we define words to correspond to Player 2 strategies. Recall that Player 2 resolves nondeterminism in  $\mathcal{G}$ . That is, when the token is in location  $l \in L$  and Player 1 selects an action  $a \in \Gamma$ , Player 2 selects an edge  $\langle l, a, l' \rangle \in \Lambda$  and the token moves to the location  $l'$ . Accordingly,  $\Sigma_I = \Lambda$ , thus a word in  $\Sigma_I^*$  is a sequence of edges. We define the specification  $\mathcal{S}$  so that a word witnesses the incorrectness of the design only if it corresponds to a *loosing path* in  $\mathcal{G}$ , where  $\langle l_0, a_1, l'_1 \rangle, \langle l_1, a_2, l'_2 \rangle, \langle l_2, a_3, l'_3 \rangle, \dots, \langle l_{m-1}, a_m, l'_m \rangle \in \Lambda^*$  is a losing path if for all  $1 \leq i \leq m-1$  we have that  $l'_i = l_i$  and  $l'_m \notin \alpha$ .

Finally, we define the components so that a correct design corresponds to a winning Player 1 strategy. For  $a \in \Gamma$ , the component  $\mathcal{B}_a \in \mathcal{L}$  can *process* every edge that is labeled with  $a$ . When reading such an edge it relinquishes control, and when reading an edge that is labeled with  $b \neq a$ , the component enters a sink which is intuitively a rejecting sink. Thus, in order to avoid processing a word  $w = t_1 \dots t_m$  that corresponds to a losing path, a design must assign control to some component  $\mathcal{B}_a$  with  $a \neq a_i$ , after reading the input prefix  $t_1 \dots t_{i-1}$ , for  $1 \leq i \leq m$ .

Formally, for  $a \in \Gamma$ , we define the box-transducer  $\mathcal{B}_a = \langle \Lambda, \{\top, \perp\}, \{c_a^0, c_a^{rej}\} \cup \mathcal{O}, c_a^0, \delta_a, \nu_a, \mathcal{O} \rangle$ , where  $\nu_a(c_a^0) = \top$ ,  $\nu_a(c_a^{rej}) = \perp$ , and  $\delta_a$  is defined as follows. Consider an edge  $\langle l, a, l' \rangle \in \Lambda$ . We define  $\delta_a(c_a^0, \langle l, a, l' \rangle) = P$ , for the observation  $P$  with  $l' \in P$ . For an edge  $\langle l, b, l' \rangle \in \Lambda$ , with  $b \neq a$ , we define  $\delta_a(c_a^0, \langle l, b, l' \rangle) = c_a^{rej}$ . The state  $c_a^{rej}$  is a sink, thus  $\delta_a(c_a^{rej}, \sigma) = c_a^{rej}$  for all  $\sigma \in \Sigma_I$ . Note that the component  $\mathcal{B}_a$  relinquishes control and outputs  $\top$  when it reads a transition labeled  $a$ . Otherwise, it gets stuck in the rejecting sink.

The specification is given by the DFA  $\mathcal{S} = \langle \Sigma_I \times \Sigma_{\mathcal{O}}, S, \delta_S, l^0, F_S \rangle$ , where  $S = L \cup \{s_{acc}\}$ , the accepting states  $F_S$  are the states that do not belong to observations in

$\alpha$ , thus  $F_S = S \setminus \bigcup_{P \in \alpha} P$ , and we describe  $\delta_S$  in the following. Consider a location  $l \in L$ . For every edge  $t = \langle l, a, l' \rangle \in \Lambda$  we define  $\delta_S(l, \langle t, \top \rangle) = l'$ . For every other letter  $\sigma \in \Sigma_I \times \Sigma_O$ , we define  $\delta_S(l, \sigma) = s_{acc}$ .

Note that a word  $w \in (\Sigma_I \times \Sigma_O)^*$  is not in  $L(S)$  if it is of the form  $w = w_1 \oplus w_2 \in (\Sigma_I \times \Sigma_O)^*$  for  $w_2 \in \top^*$  and  $w_1 \in \Sigma_I^*$  corresponds to a *loosing path* in  $\mathcal{G}$ . That is,  $w = t_1 \dots t_m$ , where  $t_1, \dots, t_m \in \Lambda$ ,  $t_1 = \langle l^0, a_1, l_1 \rangle$ , for  $1 \leq i < m$ , the target location of the transition  $t_i$  is the source location of  $t_{i+1}$ , and assuming  $t_m = \langle t_{m-1}, a_m, t_m \rangle$ , we have  $t_m \notin \alpha$ . Recall that a component  $\mathcal{B}_a \in \mathcal{L}$  relinquishes control after reading an edge that the action  $a \in \Gamma$  participates in. When reading an edge that  $a$  does not participate in,  $\mathcal{B}_a$  enters a sink and outputs  $\perp$ . Thus, in order to avoid processing the word  $w$  as in the above, a design must assign control to some component  $\mathcal{B}_a$  with  $a \neq a_i$ , after reading the input prefix  $t_1 \dots t_{i-1}$ , for  $1 \leq i \leq m$ .

We claim that Player 1 wins  $\mathcal{G}$  iff there is a correct design  $\mathcal{D}$  for  $S$  using the components in  $\mathcal{L}$ . For the first direction, consider a Player 1 winning strategy  $f_1$ . We construct a design  $\mathcal{D}$  inductively. Let  $a_1 = f_1(P_0)$ , where  $P_0 \in \mathcal{O}$  is such that  $l^0 \in P_0$ . The first component to gain control is  $\mathcal{B}_{a_1}$ . Consider a run on some word, and let  $\mathcal{B}_{a_1}, P_1, \mathcal{B}_{a_2}, P_2, \dots, \mathcal{B}_{a_m}, P_m$  be the sequence of components and corresponding exit states that the run visits. Recall that  $P_1, \dots, P_m$  are observations in  $\mathcal{O}$ . Let  $\pi = P_0, a_1, P_1, \dots, P_m$ , where  $l^0 \in P_0$ . Note that  $\pi$  might not correspond to a path in  $\mathcal{G}$ , in which case  $\mathcal{D}$  assigns control to an arbitrary component. Otherwise, assuming  $f_1(\pi) = a_{m+1}$ ,  $\mathcal{D}$  assigns control to  $\mathcal{B}_{a_{m+1}}$ .

We claim that  $\mathcal{D}$  is a correct design. Assume towards contradiction that there is a word  $w \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}}) \setminus L(S)$ . Since  $w \notin L(S)$ , its projection on  $\Sigma_O$  is in  $\top^*$  and the projection of  $w$  on  $\Sigma_I$  is a loosing path  $t_1, \dots, t_m$ . For  $1 \leq i \leq m$ , let  $t_i = \langle l_{i-1}, a_i, l_i \rangle$ , where  $l_0 = l^0$ . Let  $r$  be the run of  $\mathcal{T}_{\mathcal{L}, \mathcal{D}}$  on  $w$ . Since the projection of  $w$  on  $\Sigma_O$  is a sequence of  $\top$ 's, the sequence of components and exit states that  $r$  passes is  $\mathcal{B}_{a_1}, P_1, \dots, \mathcal{B}_{a_m}, P_m$ . Indeed, if for  $1 \leq i \leq m$  and  $b \neq a_i$ , the component  $\mathcal{B}_b \in \mathcal{L}$  is in control when  $a_i$  is read, the component outputs  $\perp$ . Let  $\pi$  be the path that corresponds to  $t_1, \dots, t_m$ , thus  $\pi = l_0, \dots, l_m$ , where  $l_0 = l^0$ . Consider the Player 2 strategy  $f_2$  that selects the edges  $t_1, \dots, t_m$ . Since by our definition of the components in  $\mathcal{L}$ , for  $1 \leq i \leq m$ , we have  $l_i \in P_i$ , we can prove by induction that  $\pi$  is a prefix of  $out(f_1, f_2)$ . Since  $l_m \notin \alpha$ , the strategy  $f_2$  is winning against  $f_1$ , which is a contradiction to the fact that it is a winning strategy, and we are done.

For the second direction, consider a correct design  $\mathcal{D}$ . We define a Player 1 strategy  $f_{\mathcal{D}}$  inductively as follows. We abuse notation and refer to the labeling function  $\nu$  of  $\mathcal{D}$  as a function from its states  $D$  to  $\mathcal{L}$ . Assume the first component that  $\mathcal{D}$  assigns control to is  $\mathcal{B}_{a_1} \in \mathcal{L}$ . For the observation  $P_0$  such that  $l^0 \in P_0$ , we define  $f_1(P_0) = a_1$ . Consider a play  $\pi = l_0, l_1, \dots, l_m$ , where  $l_0 = l^0$ . Let  $\tau = P_0, a_1, P_1, \dots, a_m, P_m$  such that for  $1 \leq i \leq m$  we have  $l_i \in P_i$ . Let  $a_1, \dots, a_m \in \Gamma$  such that, for  $1 \leq i \leq m$  we have  $\langle l_{i-1}, a_i, l_i \rangle \in \Lambda$ . Let  $\delta_{\mathcal{D}}^*(P_1 \dots P_m) = d$  and let  $\nu(d) = \mathcal{B}_{a_m} \in \mathcal{L}$ . We define  $f_{\mathcal{D}}(\tau) = a_m$ .

We claim that  $f_{\mathcal{D}}$  is a winning strategy. Assume towards contradiction that there is a Player 2 strategy  $f_2$  and a finite prefix  $\pi$  of  $out(f_{\mathcal{D}}, f_2)$  that is winning for Player 2. Let  $\tau = t_1, \dots, t_m$  be the edges that  $\pi$  traverses. Consider the run of  $\mathcal{T}_{\mathcal{L}, \mathcal{D}}$  on  $\tau$ . We can prove by induction that by our definition of  $f_{\mathcal{D}}$ , for  $1 \leq i \leq m$ , when the letter

$t_i = \langle l_{i-1}, a_i, l_i \rangle$  is read, the component  $\mathcal{B}_{a_i}$  is in control. Thus, the output of  $\mathcal{T}_{\mathcal{L}, \mathcal{D}}$  when reading the word  $\tau$  is a sequence  $w$  of  $|\tau|$   $\top$ 's. Since  $\tau$  corresponds to a losing path, we have  $\tau \oplus w \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}}) \setminus L(\mathcal{S})$ , which is a contradiction to the correctness of  $\mathcal{D}$ , and we are done.  $\blacksquare$

## 4 Libraries with Costs

Given a library and a specification, there are possibly many, in fact infinitely many, designs that are solutions to the design problem. As a trivial example, assume that  $L(\mathcal{S}) = a^*$  and that the library contains a component  $\mathcal{B}$  that traverses the letter  $a$  (that is,  $\mathcal{B}$  consists of an accepting initial state that has an  $a$ -transition to an exist state). An optimal design for  $\mathcal{S}$  uses  $\mathcal{B}$  once: it has a single state with a self loop in which  $\mathcal{B}$  is called. Other designs can use  $\mathcal{B}$  arbitrarily many numbers. When we wrote “optimal” above, we assumed that the smaller the design is, the better it is. In this section we would like to formalize the notion of optimality and add to the composition picture different costs that components in the libraries may have.

In order to capture a wide set of scenarios in practice, we associate with each component in  $\mathcal{L}$  two costs: a *construction cost* and a *quality cost*. The costs are given by the functions  $c\text{-cost}, q\text{-cost} : \mathcal{L} \rightarrow \mathbb{R}^+ \cup \{0\}$ , respectively. The construction cost of a component is the cost of adding it to the library. Thus, a design that uses a component pays its construction cost once, and (as would be the case in Section 5), when several designs use a component, they share its construction cost. The quality cost measures the performance of the component, and involves, for example, its number of states. Accordingly, a design pays the quality cost of a component every time it uses it, and the fact the component is used by other designs is not important.<sup>5</sup>

Formally, consider a library  $\mathcal{L} = \{\mathcal{B}_1, \dots, \mathcal{B}_n\}$  and a design  $\mathcal{D} = \langle [n], E, D, d^0, \delta, \nu \rangle$ . The number of times  $\mathcal{D}$  uses a component  $\mathcal{B}_i$  is  $nused(\mathcal{D}, \mathcal{B}_i) = |\{d \in D : \nu(d) = i\}|$ . The set of components that are used in  $\mathcal{D}$ , is  $used(\mathcal{D}) = \{\mathcal{B}_i : nused(\mathcal{D}, \mathcal{B}_i) \geq 1\}$ . The cost of a design is then  $cost(\mathcal{D}) = \sum_{\mathcal{B} \in used(\mathcal{D})} c\text{-cost}(\mathcal{B}) + nused(\mathcal{D}, \mathcal{B}) \cdot q\text{-cost}(\mathcal{B})$ .

We state the problem of finding the cheapest design as a decision problem. For a specification DFA  $\mathcal{S}$ , a library  $\mathcal{L}$ , and a threshold  $\mu$ , we say that an input  $\langle \mathcal{S}, \mathcal{L}, \mu \rangle$  is in BCD (standing for “bounded cost design”) iff there exists a correct design  $\mathcal{D}$  such that  $cost(\mathcal{D}) \leq \mu$ . In this section we study the BCD problem in a setting with a single user. Thus, decisions are independent of other users of the library, which, recall, may influence the construction cost.

In section 3, we reduced the design problem to the problem of the solution of a safety game. In particular, we showed how a winning strategy in the game induces a correct design. Note that while we know that safety games admits memoryless strategies, there is no guarantee that memoryless strategies are guaranteed to lead to optimal designs. We first study this point and show that, surprisingly, while memoryless strate-

<sup>5</sup> One might consider a different quality-cost model, which takes into an account the cost of *computations*. The cost of a design is then the maximal or expected cost of its computations. Such a cost model is appropriate for measures like the running time or other complexity measures. We take here a global approach, which is appropriate for measures like the number of states or security level.

gies are sufficient for obtaining an optimal design in the closed setting, this is not the case in the open setting. The source of the difference is the fact that the language of a design in the open setting may be strictly contained in the language of the specification. The approximation may enable the user to generate a design that is more complex and is still cheaper in terms of cost. This is related to the fact that over approximating the language of a DFA may result in exponentially bigger DFAs [5]. We are still able to bound the size of the cheapest design by the size of the game.

#### 4.1 On the optimality and non-optimality of memoryless strategies

Consider a closed library  $\mathcal{L}$  and a DFA  $\mathcal{S}$ . Recall that a correct design for  $\mathcal{S}$  from components in  $\mathcal{L}$  is induced by a winning strategy of Player 1 in the game  $\mathcal{G}_{\mathcal{L},\mathcal{S}}$  (see Theorem 1). If the winning strategy is not memoryless, we can trim it to a memoryless one and obtain a design whose state space is a subset of the design induced by the original strategy. Since the design has no flexibility with respect to the language of  $\mathcal{S}$ , we cannot do better. Hence the following lemma.

**Lemma 1.** *Consider a closed library  $\mathcal{L}$  and a DFA  $\mathcal{S}$ . For every  $\mu \geq 0$ , if there is a correct design  $\mathcal{D}$  with  $\text{cost}(\mathcal{D}) \leq \mu$ , then there is a correct design  $\mathcal{D}'$  induced by a memoryless strategy for Player 1 in  $\mathcal{G}_{\mathcal{L},\mathcal{S}}$  such that  $\text{cost}(\mathcal{D}') \leq \mu$ .*

**Proof:** Consider a correct design  $\mathcal{D}$  with  $\text{cost}(\mathcal{D}) \leq \mu$ . Let  $f_{\mathcal{D}}$  be a winning Player 1 strategy in the safety game  $\mathcal{G}_{\mathcal{L},\mathcal{S}}$  as constructed in the proof of Theorem 1. Note that while we know that safety games admit memoryless strategies, there is no guarantee that memoryless strategies are guaranteed to lead to minimal-cost designs. Thus,  $f_{\mathcal{D}}$  need not be a memoryless strategy. We construct a Player 1 memoryless strategy  $f'_{\mathcal{D}}$  by *trimming*  $f_{\mathcal{D}}$  in every *memoryfull* vertex. Then, we construct a design  $\mathcal{D}'$  from  $f'_{\mathcal{D}}$  and show that it costs no more than  $\mathcal{D}$ .

In order to prove the claim formally, we need a few definitions. Consider a Player 1 vertex  $v = \langle e, s \rangle \in V_1$ . We define  $\text{adj}(f_{\mathcal{D}}, v) \in V$  to be the vertices  $f_{\mathcal{D}}$  continues to from  $v$ . That is,  $u \in \text{adj}(f_{\mathcal{D}}, v)$  if there is a play  $\pi_u$  that ends in  $v$  and is an outcome of the game when Player 1 plays  $f_{\mathcal{D}}$  and Player 2 plays some strategy, and  $f_{\mathcal{D}}(\pi_u) = u$ . We refer to  $\pi_u$  as a *witness* play for  $u \in \text{adj}(f_{\mathcal{D}}, v)$ . Note that there can be many witness plays for a vertex in  $\text{adj}(f_{\mathcal{D}}, v)$ . Further note that if  $\text{adj}(f_{\mathcal{D}}, v)$  is a singleton, then  $v$  is a *memoryless* vertex with respect to  $f_{\mathcal{D}}$ . Consider a vertex  $u = \langle c_i^0, s \rangle \in \text{adj}(f_{\mathcal{D}}, v)$ . Intuitively, since we define  $f_{\mathcal{D}}$  to assign control as  $\mathcal{D}$ , and since  $f_{\mathcal{D}}$  continues to  $u$  after  $\pi_u$ , there must be a state  $d_u \in \mathcal{D}$  with  $\nu(d_u) = i$ . Thus,  $d_u$  is a *witness* state for the fact that  $f_{\mathcal{D}}(\pi_u) = u$ . Again, there can be several witness states for a vertex in  $\text{adj}(f_{\mathcal{D}}, v)$ . Formally, let  $\pi_u$  be a witness play for  $u$  and let  $w_u \in \Sigma^*$  be the word that is induced by the Player 2 choices in  $\pi_u$ . That is,  $\delta_{\mathcal{S}}^*(w_u) = s$ , where recall that  $v = \langle e, s \rangle$ . Moreover, let  $e_1, \dots, e_m$  be the exit states that are traversed in  $\pi_u$ , and let  $d_u = \delta_{\mathcal{D}}^*(e_1 \dots e_m)$ . Since  $f_{\mathcal{D}}$  assigns control to the same components as  $\mathcal{D}$ , the run of  $\mathcal{A}_{\mathcal{L},\mathcal{D}}$  on  $w_u$  ends in the state  $\langle c_i^0, d_u \rangle$ . We refer to  $w_u$  and  $d_u$  as a *witness* word and state for  $u$ , respectively.

We define  $f'_{\mathcal{D}}$  as follows. For every  $v \in V_1$ , we choose a vertex  $u \in \text{adj}(f_{\mathcal{D}}, v)$  arbitrarily and define  $f'_{\mathcal{D}}(v) = u$ . It is not hard to prove that  $f'_{\mathcal{D}}$  is a winning strategy. Let  $\mathcal{D}'$  be the design that corresponds to  $f'_{\mathcal{D}}$  and is constructed as in the proof of Theorem 1. Since  $f'_{\mathcal{D}}$  is winning,  $\mathcal{D}'$  is a correct design. We claim that  $\text{cost}(\mathcal{D}') \leq \text{cost}(\mathcal{D})$ .

Recall that the states of  $\mathcal{D}'$  are the Player 2 vertices in  $\mathcal{G}_{\mathcal{L},S}$ . Further recall that in the construction of  $\mathcal{D}'$  we skip Player 1 vertices and proceed according to  $f_{\mathcal{D}'}$ . Consider a reachable Player 2 vertex  $\langle c_i^0, s \rangle$ . We denote by  $\nu'$  the labeling function of  $\mathcal{D}'$ . Recall that  $c_i^0$  is the initial state of the component  $\mathcal{B}_i \in \mathcal{L}$  and  $\nu'(\langle c_i^0, s \rangle) = i$ . First, we claim that if a component is used in  $\mathcal{D}'$  then it is used in  $\mathcal{D}$ , thus  $used(\mathcal{D}') \subseteq used(\mathcal{D})$ . Indeed, if  $\mathcal{B}_i \in used(\mathcal{D}')$ , there is a state in  $\mathcal{D}'$  that corresponds to a reachable Player 2 vertex  $v = \langle c_i^0, s \rangle$ , for some  $s \in S$ . Let  $s_v$  be a witness state of  $v$ . Since the labeling of  $s_v$  is  $\nu(s_v) = i$ , we have  $\mathcal{B}_i \in used(\mathcal{D})$ . We conclude that the sum of construction costs that is incurred by  $\mathcal{D}'$  is at most that of  $\mathcal{D}$ .

We prove that the sum of quality costs incurred by  $\mathcal{D}'$  is at most that of  $\mathcal{D}$ . We prove that for every  $\mathcal{B}_i \in used(\mathcal{D}')$  we have  $nused(\mathcal{D}', \mathcal{B}_i) \leq nused(\mathcal{D}, \mathcal{B}_i)$ . For a reachable vertex  $v \in V_2$  let  $d_v \in D$  be an arbitrary choice of witness state of  $v$ . Let  $w_v$  be the corresponding witness word. We show that the mapping from a reachable vertex  $v \in V_2$  to  $d_v \in D$  is a one-to-one mapping. Consider  $v, u \in V_2$  with  $d_v = d_u$ . Let  $d = d_v = d_u$ . Let  $i = \nu(d)$ . Since the component-state of  $v$  and  $u$  is the initial state of  $\mathcal{B}_i$ , we have  $v = \langle c_i^0, s \rangle$  and  $u = \langle c_i^0, s' \rangle$ . To conclude the proof, we show that  $s = s'$ , thus  $v = u$ . We claim that  $L(\mathcal{S}^s) = L(\mathcal{S}^{s'})$  and since  $\mathcal{S}$  is a minimal DFA, it would follow that  $s = s'$ . Recall that  $w_v$  and  $w_u$  are the witness words for  $v$  and  $u$ , respectively, thus  $\delta_{\mathcal{S}}^*(w_v) = s$  and  $\delta_{\mathcal{S}}^*(w_u) = s'$ . Moreover, since  $w_v$  and  $w_u$  are the witness words that correspond to the witness states  $d_v = d_u = d$ , we have  $\delta_{\mathcal{L},\mathcal{D}}^*(w_v) = \delta_{\mathcal{L},\mathcal{D}}^*(w_u) = \langle c_i^0, d \rangle$ . Consider a word  $x \in \Sigma^*$ . If  $x \in L(\mathcal{S}^s)$ , then since  $\delta_{\mathcal{S}}^*(w_v) = s$ , we have  $w_v \cdot x \in L(\mathcal{S})$ . Since  $\mathcal{D}$  is a correct design,  $w_v \cdot x \in L(\mathcal{A}_{\mathcal{L},\mathcal{D}})$ . Thus,  $\delta_{\mathcal{L},\mathcal{D}}^*(w_v \cdot x)$  is an accepting state. Since  $\delta_{\mathcal{L},\mathcal{D}}^*(w_v) = \delta_{\mathcal{L},\mathcal{D}}^*(w_u)$ , we have  $w_u \cdot x \in L(\mathcal{A}_{\mathcal{L},\mathcal{D}})$  and in turn  $w_u \cdot x \in L(\mathcal{S})$ . Since  $\delta_{\mathcal{S}}^*(w_u) = s'$ , we have  $u \in L(\mathcal{S}^{s'})$ . The other direction is symmetric, and we are done. ■

While Lemma 1 seems intuitive, it does not hold in the setting of open systems. There, a design has the freedom to generate a language that is a subset of  $L(\mathcal{S})$ , as long as it stays receptive. This flexibility allows the design to generate a language that need not be related to the structure of the game  $\mathcal{G}_{\mathcal{L},S}$ , which may significantly reduce its cost. Formally, we have the following.


**Lemma 2.** *There is an open library  $\mathcal{L}$  and a family of DFAs  $\mathcal{S}_n$  such that  $\mathcal{S}_n$  has a correct design  $\mathcal{D}_n$  with cost 1 but every correct design for  $\mathcal{S}_n$  that is induced by a memoryless strategy for Player 1 in  $\mathcal{G}_{\mathcal{L},\mathcal{S}_n}$  has cost  $n$ .*

**Proof:** We define  $\mathcal{S}_n = \langle \Sigma_I \times \Sigma_O, \mathcal{S}_n, \delta_{\mathcal{S}_n}, s_0^0, F_{\mathcal{S}_n} \rangle$ , where  $\Sigma_I = \{\tilde{0}, \tilde{1}, \#\}$ ,  $\Sigma_O = \{0, 1, -\}$ , and  $\mathcal{S}_n$ ,  $\delta_{\mathcal{S}_n}$  and  $F_{\mathcal{S}_n}$  are as follows. Essentially, after reading a prefix of  $i$   $\#$ 's, for  $1 \leq i \leq n$ , the design should arrange its outputs so that the  $i$ -th and  $(n+i)$ -th letters agree (they are  $\tilde{0}$  and 0, or are  $\tilde{1}$  and 1). One method to do it is to count the number of  $\#$ 's and then check the corresponding indices. Another method is to keep track of all the first  $n$  input letters and make sure that they repeat. The key idea is that while in the second method we strengthen the specification (agreement is checked with respect to all  $i$ 's, ignoring the length of the  $\#$ -prefix), it is still receptive, which is exactly the flexibility that the open setting allows. We define  $\mathcal{S}_n$  and the library  $\mathcal{L}$  so that the structure of  $\mathcal{G}_{\mathcal{L},\mathcal{S}_n}$  supports the first method, but counting each  $\#$  has a cost of 1. Consequently, a memoryless strategy of Player 1 in  $\mathcal{G}_{\mathcal{L},\mathcal{S}_n}$  induces a design that counts,




and is therefore of cost  $n$ , whereas an optimal design follows the second method, and since it does not count the number of  $\#$ 's, its cost is only 1.

We can now describe the DFA  $\mathcal{S}_n$  in more detail. It consists of  $n$  chains, sharing an accepting sink  $s_{acc}$  and a rejecting sink  $s_{rej}$ . For  $0 \leq i \leq n-1$ , we describe the  $i$ -th chain, which is depicted in Figure 1. When describing  $\delta_{\mathcal{S}_n}$ , for ease of presentation, we sometimes omit the letter in  $\Sigma_I$  or  $\Sigma_O$  and we mean that every letter in the respective alphabet is allowed. For  $0 \leq i < n-1$ , we define  $\delta_{\mathcal{S}_n}(s_0^i, \#) = s_0^{i+1}$  and  $\delta_{\mathcal{S}_n}(s_0^n, \#) = s_0^n$ . Note that words of the form  $\#^i a_1 \bar{1} a_2 b 0$  or  $\#^i a_1 0 a_2 b 1$  are not in  $L(\mathcal{S}_n)$ , where if  $0 \leq i \leq n-1$ , then  $a_1 \in (\bar{0} + \bar{1})^i$ ,  $a_2 \in (\bar{0} + \bar{1})^{n-i-1}$ , and  $b \in (0+1)^i$ , and if  $i > n-1$  then the lengths of  $a_1$ ,  $a_2$ , and  $b$  are  $n-1$ ,  $0$ , and  $n-1$ , respectively. We require that after reading a word in  $\#^*(\bar{0} + \bar{1})^n$  there is an output of  $n$  letters in  $\{0, 1\}$ . Thus, for  $n \leq j \leq n+i+1$ , we define  $\delta_{\mathcal{S}_n}(s_j^{i,0}, -) = \delta_{\mathcal{S}_n}(s_j^{i,1}, -) = s_{rej}$ . Also,  $\mathcal{S}_n$  accepts every word that has a  $\#$  after the initial prefix of  $\#$  letters. Thus, for  $1 \leq j \leq i$ , we define  $\delta_{\mathcal{S}_n}(s_j^i, \#) = s_{acc}$ , and for  $i+1 \leq j \leq n+i+1$  and  $t \in \{0, 1\}$  we define  $\delta_{\mathcal{S}_n}(s_j^{i,t}, \#) = s_{acc}$ .



**Fig. 1.** A description of the  $i$ -th chain of the specification  $\mathcal{S}_n$ .

The library  $\mathcal{L}$  is depicted in Figure 2. The quality and construction costs of all the components is 0, except for  $\mathcal{B}_1$  which has  $q\text{-cost}(\mathcal{B}_1) = 1$  and  $c\text{-cost}(\mathcal{B}_1) = 0$ .



**Fig. 2.** The library  $\mathcal{L}$ . Exit states are square nodes and the output of a state is written in the node.

We First claim that every correct design must cost at least 1. Indeed, such a design must use  $\mathcal{B}_1$  at least once. Otherwise, the output sequence for the input word  $\#\bar{0}^{n+1}$  is a sequence in  $_*$ . Recall that we require that after reading a word in  $\#^*(\bar{0} + \bar{1})^n$  there is an output of  $n$  letters in  $\{0, 1\}$ , thus such a design is incorrect.

We describe a correct design  $\mathcal{D}_n$  with  $\text{cost}(\mathcal{D}_n) = 1$ . Intuitively, as explained above,  $\mathcal{D}_n$  does not track the number of  $\#$ 's that are read and can thus use  $\mathcal{B}_1$  only once. Instead,  $\mathcal{D}_n$  keeps track of all of the  $n$  input letters in  $\{\bar{0}, \bar{1}\}$  that follow the sequence in  $\#^*$ . Thus, after  $\bar{0}$  or  $\bar{1}$  is read,  $\mathcal{B}_2$  gains control for  $n$  consecutive times. Then, assuming the word  $u \in \{\bar{0}, \bar{1}\}^n$  is read in the first phase, control is alternated between  $\mathcal{B}_3$  and  $\mathcal{B}_4$

so that the word  $v \in \{0, 1\}^n$  is output, where  $u$  and  $v$  represent the same vector. Note that  $\mathcal{D}_n$  uses the components  $\mathcal{B}_2$ ,  $\mathcal{B}_3$ , and  $\mathcal{B}_4$  exponentially many times in  $n$ .

Formally, we define  $\mathcal{D}_n = \langle \Sigma_I, \Sigma_O, D_n, \delta_{\mathcal{D}_n}, d^0, \nu \rangle$ . We define  $\nu(d^0) = 1$  and  $\delta_{\mathcal{D}_n}(d^0, e_{\#}^1) = d^0$ . When  $\mathcal{B}_1$  relinquishes control through  $e_1^1$  or  $e_0^1$  control is passed to  $\mathcal{B}_2$  for  $n$  consecutive times. During this phase,  $\mathcal{D}_n$  remembers the vector in  $\{\tilde{0}, \tilde{1}\}^n$  that is read, thus  $\mathcal{D}_n$  has  $2^{\Omega(n)}$  states. Next, the components  $\mathcal{B}_3$  and  $\mathcal{B}_4$  alternate control for  $n$  times as we describe in the following. Consider an input word  $w = \#^i \cdot a_1 \dots a_n b$ , where  $a_1, \dots, a_n \in \{\tilde{0}, \tilde{1}\}^n$  and  $b \in \{\tilde{0}, \tilde{1}\}^m$  for  $0 \leq m < n-1$ . After reading  $w$  either  $\mathcal{B}_3$  or  $\mathcal{B}_4$  relinquish control. If  $a_{m+1} = 0$ , then the next component to gain control is  $\mathcal{B}_3$  and otherwise it is  $\mathcal{B}_4$ . After  $n$  alternations of control  $\mathcal{B}_4$  is assigned control indefinitely.

Since the initial state of  $\mathcal{D}_n$  is the only state that is labeled with  $\mathcal{B}_1$ , we have  $nused(\mathcal{B}_1) = 1$ , thus  $cost(\mathcal{D}_n) = 1$ . We claim that  $\mathcal{D}_n$  is a correct design. Indeed, note that  $L(\mathcal{T}_{\mathcal{L}, \mathcal{D}_n})$  consists of two types of words. The first type are prefixes of words in  $\#^*ab0^*$ , for  $a \in \{\tilde{0}, \tilde{1}\}^n$  and  $b \in \{0, 1\}^n$  that represent the same vector, thus for  $1 \leq i \leq n$ , if  $a_i = \tilde{0}$ , then  $b_i = 0$  and if  $a_i = \tilde{1}$ , then  $b_i = 1$ . The second type are words that have a prefix in  $\#^*(\tilde{0} + \tilde{1})^+\#$ . Clearly,  $L(\mathcal{T}_{\mathcal{L}, \mathcal{D}_n}) \subseteq L(\mathcal{S}_n)$ , and we are done. ■

## 4.2 Solving the BCD problem

**Theorem 4.** *The BCD problem is NP-complete for closed designs.*

**Proof:** Consider an input  $\langle \mathcal{S}, \mathcal{L}, \mu \rangle$  to the BCD problem. By Lemma 1, we can restrict the search for a correct design  $\mathcal{D}$  with  $cost(\mathcal{D}) \leq \mu$  to these induced by a memoryless strategy for Player 1 in  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$ . By the definition of the game  $\mathcal{G}_{\mathcal{L}, \mathcal{S}}$ , such a design has at most  $|\mathcal{C}_0 \times \mathcal{S}|$  states. Since checking if a design is correct and calculating its cost can be done in polynomial time, membership in NP follows.

For the lower bound, we describe a reduction from SET-COVER. Consider an input  $\langle U, T, \mu \rangle$  to SET-COVER, where  $U = \{1, \dots, l\}$  is a universe,  $T = \{T_1, \dots, T_m\}$  is a set of subsets over  $U$ , i.e.,  $T_i \subseteq U$ , for  $1 \leq i \leq m$ , and  $\mu \in \mathbb{N}$ . We construct a closed-library  $\mathcal{L}$  and a DFA  $\mathcal{S}$  such that there is a design  $\mathcal{D}$  that costs at most  $\mu$  iff there is a set-cover of size at most  $\mu$ .

We construct the DFA  $\mathcal{S}$  over the alphabet  $U$  as a chain of  $l+1$  states that accepts only the word  $12\dots l$ . There are  $m+2$  components in  $\mathcal{L}$ . For  $1 \leq i \leq m$ , the component  $\mathcal{B}_i$  corresponds to the set  $T_i$ . Its initial state is  $c_i^0$  and it has an exit state  $e_j^i$  for every element  $j \in T_i$ . For  $j \in T_i$  there is a transition  $\delta_i(c_i^0, j) = e_j^i$ . The last two components are  $\mathcal{B}_{acc}$  and  $\mathcal{B}_{rej}$ . For consistency we also refer to these components as  $\mathcal{B}_{m+1}$  and  $\mathcal{B}_{m+2}$ , respectively. They consist of a single initial non-exit state that is accepting in  $\mathcal{B}_{acc}$  and rejecting in  $\mathcal{B}_{rej}$ . For  $\mathcal{B}_i \in \mathcal{L}$ , we define  $c-cost(\mathcal{B}_i) = 1$  and we define  $c-cost(\mathcal{B}_{acc}) = c-cost(\mathcal{B}_{rej}) = 0$ . Finally, the quality cost of the components is 0, thus we define  $q-cost \equiv 0$ .

We claim that there is a set cover of size at most  $\mu$  iff there is a design of cost at most  $\mu$ . For the first direction, consider a set cover  $T' \subseteq T$  with  $|T'| \leq \mu$ . For  $j \in U$ , let  $cover(j)$  be the index such that  $j \in T_{cover(j)}$  and  $T_{cover(j)} \in T'$ . We construct a design  $\mathcal{D}$  with  $cost(\mathcal{D}) \leq \mu$ . We define  $\mathcal{D}$  so that when  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  reads the word  $1\dots l$ , the component that gets  $j$ -th control, for  $1 \leq j \leq l$ , is  $\mathcal{B}_{cover(j)}$ . If the  $j$ -th input letter

is  $j' \neq j$ , then the next component to gain control is  $\mathcal{B}_{rej}$ . Finally, the component that gains  $l + 1$ -th control is  $\mathcal{B}_{acc}$ .

Formally,  $\mathcal{D} = \langle \mathcal{E}, [m + 2], D, \delta_{\mathcal{D}}, d_1, \nu \rangle$ , where  $\{d_1, \dots, d_l, d_{acc}, d_{rej}\}$  and we define  $\nu$  and  $\delta_{\mathcal{D}}$  as follows. For  $1 \leq j \leq l$ , we define  $\nu(d_j) = \text{cover}(j)$ . For  $1 \leq i \leq m$ , recall that the component  $\mathcal{B}_i$  has an exit state  $e_j^i$  for every  $j \in T_i$ . For  $j < l$  and  $1 \leq i \leq m$ , we define  $\delta_{\mathcal{D}}(d_j, e_j^i) = d_{j+1}$  and we define  $\delta_{\mathcal{D}}(d_l, e_l^i) = d_{acc}$ . For  $j \neq i$ , we define  $\delta_{\mathcal{D}}(d_j, e_j^i) = d_{rej}$ . It is not hard to see that  $L(\mathcal{A}_{\mathcal{L}, \mathcal{D}}) = \{12 \dots l\}$ , and thus  $\mathcal{D}$  is a correct design. The components that  $\mathcal{D}$  uses are the ones that correspond to the sets in  $T'$ , thus  $\mathcal{D}$  uses at most  $\mu$  components with construction cost 1, and its cost is at most  $\mu$ .

For the other direction, assume there is a correct design  $\mathcal{D}$  with  $\text{cost}(\mathcal{D}) \leq \mu$ . Consider the collection  $T' \subseteq T$  of sets that correspond to the components  $\{\nu(d) : d \in D \text{ and } \nu(d) \notin \{\mathcal{B}_{acc}, \mathcal{B}_{rej}\}\}$ . We claim that  $T'$  is a set cover with at most  $\mu$  sets. Since  $\text{cost}(\mathcal{D}) \leq \mu$ , it uses at most  $\mu$  components that correspond to sets in  $T$ , thus  $|T'| \leq \mu$ . We are left to show that  $T'$  is a set cover. Consider the run  $r$  of  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  on the word  $1 \dots l$ . Since  $\mathcal{D}$  is correct,  $r$  is accepting. Specifically it does not get stuck. Since the components we described above relinquish control after reading a single letter, the control is passed  $l + 1$  times during  $r$ . Let  $\mathcal{B}_{i_1}, \dots, \mathcal{B}_{i_l}, \mathcal{B}_{i_{l+1}}$  be the sequence of components that gain control. Consider  $j \in U$ . We claim that  $j$  is covered in  $T'$ . Since  $r$  does not get stuck, there is a transition labeled  $j$  in the component  $\mathcal{B}_{i_j}$ , the  $j$ -th component to gain control in  $r$ . Thus,  $j \in T_{i_j}$ , the set in  $T$  that corresponds to  $\mathcal{B}_{i_j}$ . Since  $T_{i_j} \in T'$ ,  $j$  is covered, and we are done. ■

*Remark 1.* Note that a different attempt to reduce SET-COVER to the BCD problem would be to define the components as in the proof of Theorem 4 and define  $\mathcal{S}$  so that it accepts one-letter words for each element in  $U$ . However, this attempt fails since, intuitively, a design does not know which letter is going to be read. Even if there is a set cover  $T' \subseteq T$  and control is assigned to the component  $\mathcal{B}_i$  where  $T_i \in T'$ , a letter  $j \in U \setminus T_i$  can be read. Thus, the fact that we use a one-word specification allows a design to expect the next letter that should be read.

We turn to study the open setting, which is significantly harder than the closed one. For the upper bound, we first show that while we cannot restrict attention to designs induced by memoryless strategies, we can still bound the size of optimal designs:

**Theorem 5.** *For an open library  $\mathcal{L}$  with  $\ell$  components and a specification  $\mathcal{S}$  with  $n$  states, a cheapest correct design  $\mathcal{D}$  has at most  $\binom{n}{n/2} \cdot \ell$  states.*

**Proof:** Given  $\mathcal{S}$  and  $\mathcal{L}$ , assume towards contradiction that the cheapest smallest design  $\mathcal{D}$  for  $\mathcal{S}$  using the components in  $\mathcal{L}$  has more than  $\binom{n}{n/2} \cdot \ell$  states.

Consider a word  $w \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}})$ . Let  $\mathcal{B}_{i_1}, \dots, \mathcal{B}_{i_m} \in \mathcal{L}$  be the components that are traversed in the run  $r$  of  $\mathcal{T}_{\mathcal{L}, \mathcal{D}}$  that induces  $w$ . Let  $w = w_1 \dots w_m$ , where, for  $1 \leq j \leq m$ , the word  $w_j$  is induced in the component  $\mathcal{B}_{i_j}$ . We say that  $w$  is suffix-less if  $w_m = \epsilon$ , thus  $r$  ends in the initial state of the last component to gain control. We denote by  $\pi_w(\mathcal{D}) = e_{i_1}, \dots, e_{i_{m-1}} \in \mathcal{E}^*$  the sequence of exit states that  $r$  visits.

For a state  $d \in D$ , we define the set  $S_d \subseteq S$  so that  $s \in S_d$  iff there exists a suffix-less word  $w \in (\Sigma_I \times \Sigma_O)^*$  such that  $\delta_{\mathcal{S}}^*(w) = s$  and  $\delta_{\mathcal{D}}^*(\pi_w(\mathcal{D})) = d$ . Since  $\mathcal{D}$

has more than  $\binom{n}{n/2} \cdot \ell$  states, there is a component  $\mathcal{B}_i \in \mathcal{L}$  such that the set  $D' \subseteq \mathcal{L}$  of states that are labeled with  $\mathcal{B}_i$  is larger than  $\binom{n}{n/2}$ . Thus, there must be two states  $d, d' \in D'$  that have  $S_{d'} \subseteq S_d$ . Note that  $\nu(d) = \nu(d') = i$ .

We construct a new design  $\mathcal{D}'$  by merging  $d'$  into  $d$ . Formally, we define  $D' = D \setminus \{d'\}$ . If  $d'$  is the initial state of  $\mathcal{D}$ , then we define  $d'^0 = d$  and otherwise we define  $d'^0 = d^0$ . For  $t \in D'$  and  $e \in \mathcal{E}$ , if  $\delta_{\mathcal{D}}(t, e) = d'$ , then we define  $\delta_{\mathcal{D}'}(t, e) = d$ , and otherwise we define  $\delta_{\mathcal{D}'}(t, e) = \delta_{\mathcal{D}}(t, e)$ . Finally, for every  $t \in D'$  we define  $\nu'(t) = \nu(t)$ .

Clearly, for every component  $\mathcal{B} \in \mathcal{L}$  we have  $nused(\mathcal{D}, \mathcal{B}) \geq nused(\mathcal{D}', \mathcal{B})$ , thus  $cost(\mathcal{D}) \geq cost(\mathcal{D}')$ . Moreover,  $\mathcal{D}'$  has less states than  $\mathcal{D}$ . Thus, in order complete the contradiction and conclude the proof of the theorem, we prove the following.

*Claim.* We claim that  $\mathcal{D}'$  is a correct design, thus  $L(\mathcal{T}_{\mathcal{L}, \mathcal{D}'}) \subseteq L(\mathcal{S})$ .

In order to prove the claim we prove the following.

*Claim.* For every suffix-less (w.r.t  $\mathcal{D}'$ ) word  $x \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}'})$  there is a suffix-less (w.r.t  $\mathcal{D}$ ) word  $y \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}})$  such that  $\delta_{\mathcal{S}}^*(x) = \delta_{\mathcal{S}}^*(y)$  and  $\delta_{\mathcal{D}}^*(\pi_y(\mathcal{D})) = \delta_{\mathcal{D}'}^*(\pi_x(\mathcal{D}'))$ .

Before proving the correctness of Claim 4.2 we show that it implies the correctness of Claim 4.2. Assume towards contradiction that there is a word  $w \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}'}) \setminus L(\mathcal{S})$ . Let  $r$  be the run of  $\mathcal{T}_{\mathcal{L}, \mathcal{D}'}$  that induces  $w$  and let  $\mathcal{B}_i \in \mathcal{L}$  be the last component to gain control in  $r$ . Let  $w_0, \dots, w_m$  be the partition of  $w$  with respect to  $\mathcal{D}'$  and let  $x = w_0 \cdot \dots \cdot w_{m-1}$ . That is,  $x$  is the longest suffix-less prefix of  $w$ . Note that  $\pi_w(\mathcal{D}') = \pi_x(\mathcal{D}')$ . By Claim 4.2 there is a suffix-less (w.r.t  $\mathcal{D}$ ) word  $y \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}})$  such that  $\delta_{\mathcal{S}}^*(x) = \delta_{\mathcal{S}}^*(y)$  and  $\delta_{\mathcal{D}}^*(\pi_y(\mathcal{D})) = \delta_{\mathcal{D}'}^*(\pi_x(\mathcal{D}'))$ . Since  $\delta_{\mathcal{D}}^*(\pi_y(\mathcal{D})) = \delta_{\mathcal{D}'}^*(\pi_x(\mathcal{D}'))$ , the last component to gain control in the two runs is  $\mathcal{B}_i$ . Since both  $x$  and  $y$  are suffix-free, the runs that induce them end in the initial state of  $\mathcal{B}_i$ , thus  $\delta_{\mathcal{L}, \mathcal{D}}^*(y) = \delta_{\mathcal{L}, \mathcal{D}'}^*(x)$ . Recall that  $w = x \cdot w_n \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}'})$ . Thus,  $y \cdot w_n \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}})$ . Since  $\delta_{\mathcal{S}}^*(y) = \delta_{\mathcal{S}}^*(x)$  and  $w \notin L(\mathcal{S})$ , we have  $y \cdot w_n \notin L(\mathcal{S})$ . Thus,  $y \cdot w_n \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}}) \setminus L(\mathcal{S})$ , and we reach a contradiction to the correctness of  $\mathcal{D}$ , and we are done.

We continue to prove Claim 4.2. Consider a suffix-less (w.r.t  $\mathcal{D}'$ ) word  $x \in L(\mathcal{T}_{\mathcal{L}, \mathcal{D}'})$  and let  $r$  be the run of  $\mathcal{D}'$  on  $\pi_x(\mathcal{D}')$ , which are the exit states that  $\mathcal{T}_{\mathcal{L}, \mathcal{D}'}$  visits when inducing  $x$ . The proof is by induction on the number of visits of  $r$  to  $d \in D'$ . For the base case,  $r$  does not visit  $d$ . Thus,  $r$  is a run of  $\mathcal{D}$  on  $\pi_x(\mathcal{D}')$ , and we define  $y = x$ .

Assume the claim is correct for runs with  $m$  visits to  $d$  and we prove for runs with  $m + 1$  visits. Let  $\pi_x(\mathcal{D}') = \pi_{x_1}(\mathcal{D}') \cdot e \cdot \pi_{x_3}(\mathcal{D}')$  such that  $\delta_{\mathcal{D}'}^*(\pi_{x_1}(\mathcal{D}') \cdot e) = d$  is the last visit of  $r$  to  $d$ . Let  $x = x_1 \cdot x_2 \cdot x_3$ , where  $x_1$  and  $x_1 \cdot x_2$  are suffix-less. That is, assuming  $t = \delta_{\mathcal{D}'}^*(\pi_{x_1}(\mathcal{D}'))$  and  $\nu'(t) = i$ , then  $x_2$  is the sub word of  $x$  that is induced by the component  $\mathcal{B}_i$ , thus  $\delta_i^*(x_2) = e \in \mathcal{E}$ . We distinguish between two cases. In the first case  $x_2 = \epsilon$ , thus  $x_1 = \epsilon$  and  $d'^0 = d$ , and we define  $y = x$  as in the base case. In the second case,  $x_2 \neq \epsilon$ . By the induction hypothesis, there is a suffix-less (w.r.t  $\mathcal{D}$ ) word  $y_1 \in (\Sigma_I \times \Sigma_O)^*$  such that  $\delta_{\mathcal{D}}^*(\pi_{y_1}(\mathcal{D})) = \delta_{\mathcal{D}'}^*(\pi_{x_1}(\mathcal{D}')) = t \in D'$  and  $\delta_{\mathcal{S}}^*(y_1) = \delta_{\mathcal{S}}^*(x_1)$ . If  $\delta_{\mathcal{D}}(t, e) = d$ , then  $\delta_{\mathcal{D}'}(t, e) = d$  and  $y = y_1 \cdot x_2 \cdot x_3$  clearly meets the requirements of the claim. Assume  $\delta_{\mathcal{D}}(t, e) = d'$ . Note that  $y_1 \cdot x_2$  is a suffix-less (w.r.t  $\mathcal{D}$ ) word with  $\delta_{\mathcal{D}}^*(\pi_{y_1 \cdot x_2}(\mathcal{D})) = d'$ . Let  $\delta_{\mathcal{S}}^*(y_1 \cdot x_2) = s$ . Since  $S_{d'} \subseteq S_d$ , there is a suffix-less (w.r.t  $\mathcal{D}$ ) word  $z \in (\Sigma_I \times \Sigma_O)^*$  such that  $\delta_{\mathcal{D}}(\pi_z(\mathcal{D})) = d$  and  $\delta_{\mathcal{S}}^*(z) = s$ .

We define  $y = z \cdot x_3$ , which clearly meets the requirements of the claim, and we are done.  $\blacksquare$

Before we turn to the lower bound, we argue that the exponential blow-up proven in Theorem 5 cannot be avoided:

**Theorem 6.** *For every  $n \geq 1$ , there is an open library  $\mathcal{L}$  and specification  $\mathcal{S}_n$  such that the size of  $\mathcal{L}$  is constant, the size of  $\mathcal{S}_n$  is  $O(n^2)$ , and every cheapest correct design for  $\mathcal{S}_n$  that uses components from  $\mathcal{L}$  has at least  $2^n$  states.*

**Proof:** Consider the specification  $\mathcal{S}_n$  and library  $\mathcal{L}$  that are described in Lemma 2. We claim that every correct design that costs 1 cannot count  $\#$ 's and should thus remember vectors in  $\{\tilde{0}, \tilde{1}\}^n$ . Assume towards contradiction that there is a correct design  $\mathcal{D}$  with  $\text{cost}(\mathcal{D}) = 1$  and  $\mathcal{D}$  has less than  $2^n$  states. Thus,  $n_{\text{used}}(\mathcal{D}, \mathcal{B}_1) = 1$ . Since  $\mathcal{D}$  must assign initial control to  $\mathcal{B}_1$ , its initial state  $d^0$  is labeled with 1. We claim that if  $\mathcal{B}_1$  relinquishes control after reading  $\#$ , it is assigned control again. Indeed, if one of the other components gains control, for the input word  $\#\tilde{0}^{n+1}$  the output sequence is a sequence in  $_*$ . Recall that we require that after reading a word in  $\#^*(\tilde{0} + \tilde{1})^n$  there is an output of  $n$  letters in  $\{0, 1\}$ , thus we reach a contradiction to the correctness of  $\mathcal{D}$ . We conclude that the initial state has a  $e_{\#}^1$ -labeled self loop, thus  $\delta_{\mathcal{D}}(d^0, e_{\#}^1) = d^0$ .

Since  $\mathcal{D}$  has less than  $2^n$  states and the components  $\mathcal{B}_2, \mathcal{B}_3$ , and  $\mathcal{B}_4$  relinquish control after reading a letter in  $\{\tilde{0}, \tilde{1}\}$ , there are two words  $a, b \in \{\tilde{0}, \tilde{1}\}^n$  such that  $a \neq b$  and  $\delta_{\mathcal{L}, \mathcal{D}}^*(a) = \delta_{\mathcal{L}, \mathcal{D}}^*(b)$ . Let  $0 \leq i \leq n - 1$  such that, WLog,  $1 = a_i \neq b_i = 0$ . By the above, for the initial state  $q_{\mathcal{L}, \mathcal{D}}^0$  of  $\mathcal{T}_{\mathcal{L}, \mathcal{D}}$  we have  $\delta_{\mathcal{L}, \mathcal{D}}(q_{\mathcal{L}, \mathcal{D}}^0, \#) = q_{\mathcal{L}, \mathcal{D}}^0$ . Thus,  $\delta_{\mathcal{L}, \mathcal{D}}^*(\#^i a \tilde{0}^i) = \delta_{\mathcal{L}, \mathcal{D}}^*(\#^i b \tilde{0}^i)$ . Recall that  $\mathcal{S}_n$  requires that either 0 or 1 is output after these words are read, thus either  $\mathcal{B}_3$  or  $\mathcal{B}_4$  gain control. In the first case, the output letter is 0 and the input word  $\#^i a \tilde{0}^{i+1}$  produces an output that violates the specification, and in the second case 1 is output and the input word  $\#^i b \tilde{0}^{i+1}$  produces an output that violates the specification, and we are done.  $\blacksquare$

**Theorem 7.** *The BCD problem for open libraries is NEXPTIME-complete.*

**Proof:** Membership in NEXPTIME follows from Theorem 5 and the fact we can check the correctness of a design and calculate its cost in polynomial time.

For the lower bound, we describe a reduction from *EXP-TILING*. Consider an input to *EXP-TILING*  $\langle T, V, H, n \rangle$ , where  $T = \{t_1, \dots, t_m\}$  is a set of tiles,  $V, H \subseteq T \times T$  are *vertical* and *horizontal* relations, respectively, and  $n \in \mathbb{N}$  is an index. We say that  $\langle T, V, H, n \rangle \in \text{EXP-TILING}$  if it is possible to fill a  $2^n \times 2^n$  square with the tiles in  $T$  that respects the two relations. Formally,  $\langle T, V, H, n \rangle \in \text{EXP-TILING}$  if there is a function  $f : 2^n \times 2^n \rightarrow T$  such that for  $a, b \in 2^n$ , if  $a < 2^n$  then  $V(f(a, b), f(a+1, b))$ , and if  $b < 2^n$ , then  $H(f(a, b), f(a, b+1))$ .

Given an input  $\langle T, V, H, n \rangle$ , we construct an input  $\langle \mathcal{L}, \mathcal{S}, k \rangle$  to the open-BCD problem such that there is an exponential tiling iff there is a correct design  $\mathcal{D}$  with  $\text{cost}(\mathcal{D}) \leq 2^{2n+1} + 1$ . The idea behind the reduction is similar to that of Lemma 2. We define  $\Sigma_I = \{\tilde{0}, \tilde{1}, \#, c, v, h, -\}$  and  $\Sigma_O = \{0, 1, -\} \cup T$ . For  $x \in \{0, 1\}^n$ , we use  $\tilde{x}$  to refer to the  $\{\tilde{0}, \tilde{1}\}$  copy of  $x$ . The library  $\mathcal{L}$  has the same components as in Lemma 2 with an additional *tile component*  $\mathcal{B}_t$  for every  $t \in T$ . The component  $\mathcal{B}_t$  outputs  $t$  in its initial state, and when reading  $c, v$ , or  $h$ , it relinquishes control. When reading every

other letter, it enters an accepting sink. The construction costs of the components in  $\mathcal{L}$  is 0. We define  $q\text{-cost}(\mathcal{B}_1) = 2^{2^n} + 1$ , and  $q\text{-cost}(\mathcal{B}_t) = 1$  for all  $t \in T$ . The other components' quality cost is 0.

Consider a correct design  $\mathcal{D}$  with  $\text{cost}(\mathcal{D}) \leq 2^{2^{n+1}} + 1$ . We define  $\mathcal{S}$  so that a correct design must use  $\mathcal{B}_1$  at least once, thus  $\mathcal{D}$  uses it exactly once. Intuitively,  $a \cdot b$ , for  $a, b \in \{0, 1\}^n$ , can be thought of as two coordinates in a  $2^n \times 2^n$  square. We define  $\mathcal{S}$  so that after reading the word  $\tilde{a} \cdot \tilde{b} \in \{\tilde{0}, \tilde{1}\}^{2^n}$ , a component is output, which can be thought of as the tile in the  $(a, b)$  coordinate in the square. The next letter that can be read is either  $c, v$ , or  $h$ . Then,  $\mathcal{S}$  enforces that the output is  $a \cdot b$ ,  $(a+1) \cdot b$ , and  $a \cdot (b+1)$ , respectively. Thus, we show that  $\mathcal{D}$  uses exactly  $2^{2^n}$  tile components and the tiling that it induces is legal.

The other direction is more involved. Assume there is a correct design  $\mathcal{D}$  with  $\text{cost}(\mathcal{D}) \leq 2^{2^{n+1}} + 1$ . We define  $\mathcal{S}$  so that it guarantees that  $\mathcal{B}_1$  is used at least once. Thus,  $\mathcal{D}$  uses  $\mathcal{B}_1$  exactly once and uses at most  $2^{2^n}$  tile components. We describe the specification  $\mathcal{S}$  as a union of the languages of five DFAs. The first DFA  $\mathcal{S}_c$  is similar to the specification in Theorem 6. The differences are that there are  $2n$  chains and after the sequence of  $2n$  input letters, a letter  $\langle t, c \rangle$ , for  $t \in T$ , must be read. Thus, it guarantees that when  $\mathcal{D}$  reads  $\tilde{a}\tilde{b}c_{-2^n}$ , for  $\tilde{a}, \tilde{b} \in \{\tilde{0}, \tilde{1}\}^{2^n}$ , it outputs a word  $_{-2^n}tab$ , where  $t \in T$  and recall that  $a, b \in \{0, 1\}^n$  represent the same vectors as  $\tilde{a}$  and  $\tilde{b}$ , respectively. We define a tiling  $f$  by  $f(a, b) = \nu(\delta_{\mathcal{L}, \mathcal{D}}(\tilde{a}\tilde{b}))$  for  $a, b \in \{0, 1\}^n$ . The next DFA from which  $\mathcal{S}$  is devised is  $\mathcal{S}_v$ . It guarantees that when  $\mathcal{D}$  reads the word  $\tilde{a}\tilde{b}vc_{-2^n}$  it outputs  $_{-2^n}t_it_j(a+1)b$  for  $t_i, t_j \in T$  with  $V(t_i, t_j)$ . Note that since  $\mathcal{D}$  uses at most  $2^{2^n}$  tile components we have  $f(a+1, b) = t_j$ , thus  $f$  respects  $V$ . The DFA  $\mathcal{S}_h$  is defined similarly to  $\mathcal{S}_v$ . Finally, in order to deal with the tiles in the top row and left column, we have  $\mathcal{S}_{v,1}$  and  $\mathcal{S}_{h,1}$  that accepts every word that starts with a prefix in  $\#^* \tilde{1}^n (\tilde{0} + \tilde{1})^n v$  and  $\#^* (\tilde{0} + \tilde{1})^n \tilde{1}^n h$ , respectively. ■

## 5 Libraries with Costs and Multiple Users

In this section we study the setting in which several designers, each with his own specification, use the library. The construction cost of a component is now shared by the designers that use it, with the share being proportional to the number of times the component is used. For example, if  $c\text{-cost}(\mathcal{B}) = 8$  and there are two designers, one that uses  $\mathcal{B}$  once and a second that uses  $\mathcal{B}$  three times, then the construction costs of  $\mathcal{B}$  of the two designers are 2 and 6, respectively. The quality cost of a component is not shared. Thus, the cost a designer pays for a design depends on the choices of the other users and he has an incentive to share the construction costs of components with other designers. We model this setting as a multi-player game, which we dub *component library games* (CLGs, for short). The game can be thought of as a one-round game in which each player (user) selects a design that is correct according to his specification. In this section we focus on closed designs.


Formally, a CLG is a tuple  $\langle \mathcal{L}, \mathcal{S}_1, \dots, \mathcal{S}_k \rangle$ , where  $\mathcal{L}$  is a closed component library and, for  $1 \leq i \leq k$ , the DFA  $\mathcal{S}_i$  is a specification for Player  $i$ . A strategy of Player  $i$  is a design that is correct with respect to  $\mathcal{S}_i$ . We refer to a choice of designs for all the players as a *strategy profile*. Consider a profile  $P = \langle \mathcal{D}_1, \dots, \mathcal{D}_k \rangle$  and a component  $\mathcal{B} \in \mathcal{L}$ . The

construction cost of  $\mathcal{B}$  is split proportionally between the players that use it. Formally, for  $1 \leq i \leq k$ , recall that we use  $nused(\mathcal{B}, \mathcal{D}_i)$  to denote the number of times  $\mathcal{D}_i$  uses  $\mathcal{B}$ . For a profile  $P$ , let  $nused(\mathcal{B}, P)$  denote the number of times  $\mathcal{B}$  is used by all the designs in  $P$ . Thus,  $nused(\mathcal{B}, P) = \sum_{1 \leq i \leq k} nused(\mathcal{B}, \mathcal{D}_i)$ . Then, the construction cost that Player  $i$  pays in  $P$  for  $\mathcal{B}$  is  $c\text{-cost}_i(P, \mathcal{B}) = c\text{-cost}(\mathcal{B}) \cdot \frac{nused(\mathcal{B}, \mathcal{D}_i)}{nused(\mathcal{B}, P)}$ . Since the quality costs of the components is not shared, it is calculated as in Section 4. Thus, the cost Player  $i$  pays in profile  $P$ , denoted  $cost_i(P)$  is  $\sum_{\mathcal{B} \in \mathcal{L}} c\text{-cost}_i(P, \mathcal{B}) + nused(\mathcal{B}, \mathcal{D}_i) \cdot q\text{-cost}(\mathcal{D}_i)$ . We define the cost of a profile  $P$ , denoted  $cost(P)$ , as  $\sum_{i \in [k]} cost_i(P)$ .

For a profile  $P$  and a correct design  $\mathcal{D}$  for Player  $i$ , let  $P[i \leftarrow \mathcal{D}]$  denote the profile obtained from  $P$  by replacing the choice of design of Player  $i$  by  $\mathcal{D}$ . A profile  $P$  is a *Nash equilibrium* (NE) if no Player  $i$  can benefit by unilaterally deviating from his choice in  $P$  to a different design; i.e., for every Player  $i$  and every correct design  $\mathcal{D}$  with respect to  $\mathcal{S}_i$ , it holds that  $cost_i(P[i \leftarrow \mathcal{D}]) \geq cost_i(P)$ .

**Theorem 8.** *There is a CLG with no NE.*

**Proof:** We adapt the example for multiset cost-sharing games from [6] to CLGs. Consider the two-player CLG over the alphabet  $\Sigma = \{a, b, c\}$  in which Player 1 and 2's specifications are (the single word) languages  $\{ab\}$  and  $\{c\}$ , respectively. The library is depicted in Figure 3, where the quality costs of all components is 0,  $c\text{-cost}(\mathcal{B}_1) = 12$ ,  $c\text{-cost}(\mathcal{B}_2) = 5$ ,  $c\text{-cost}(\mathcal{B}_3) = 1$ , and  $c\text{-cost}(\mathcal{B}_4) = c\text{-cost}(\mathcal{B}_5) = 0$ . Both players have two correct designs. For Player 1, the first design uses  $\mathcal{B}_1$  twice and the second design uses  $\mathcal{B}_1$  once and  $\mathcal{B}_2$  once. There are also uses of  $\mathcal{B}_4$  and  $\mathcal{B}_5$ , but since they can be used for free, we do not include them in the calculations. For Player 2, the first design uses  $\mathcal{B}_2$  once, and the second design uses  $\mathcal{B}_1$  once. The table in Figure 3 shows the players' costs in the four possible CLG's profiles, and indeed none of the profiles is a NE. ■



**Fig. 3.** The library of the CLG with no NE, and the costs of the players in its profiles.

We study computational problems for CLGs. The most basic problem is the *best-response problem* (BR problem, for short). Given a profile  $P$  and  $i \in [k]$ , find the cheapest correct design for Player  $i$  with respect to the other players' choices in  $P$ . Apart from its practical importance, it is an important ingredient in the solutions to the other problems we study. The next problem we study is finding the *social optimum* (SO, for short), namely the profile that minimizes the total cost of all players; thus the one obtained when the players obey some centralized authority. For both the BR and SO problems, we study the decision (rather than search) variants, where the input includes a threshold  $\mu$ . Finally, since CLGs are not guaranteed to have a NE, we study the problem of deciding whether a given CLG has a NE. We term this problem  $\exists$ NE.

**Definition 1.** *We define the decision problems formally as follows. Let  $\mathcal{G}$  be a CLG.*

**BR** An input  $\langle \mathcal{G}, P, i, \mu \rangle$  is in BR, where  $P$  is a profile,  $i \in [k]$ , and  $\mu \in \mathbb{R}$ , iff there is a design  $\mathcal{D}_i$  that is correct with respect to  $\mathcal{S}_i$  such that  $\text{cost}_i(P[i \leftarrow \mathcal{D}_i]) \leq \mu$ .

**SO** An input  $\langle \mathcal{G}, \mu \rangle$  is in SO, where  $\mu \in \mathbb{R}$ , iff there is a profile  $P$  such that  $\text{cost}(P) \leq \mu$ .

**$\exists$ NE** An input  $\langle \mathcal{G} \rangle$  is in  $\exists$ NE iff there is a NE profile in  $\mathcal{G}$ .

Note that the BCD problem studied in Section 4 is a special case of BRP when there is only one player. Also, in a setting with a single player, the SO and BR problems coincide, thus the lower bound of Theorem 4 applies to them. In Lemma 1 we showed that if there is a correct design  $\mathcal{D}$  with  $\text{cost}(\mathcal{D}) \leq \mu$ , then there is also a correct design  $\mathcal{D}'$ , based on a memoryless strategy and hence having polynomially many states, such that for every component  $\mathcal{B}$ , we have  $nused(\mathcal{D}', \mathcal{B}) \leq nused(\mathcal{D}, \mathcal{B})$ . The arguments there apply in the more general case of CLGs. Thus, we have the following.

**Theorem 9.** *The BR and SO problems are NP-complete.*

We continue to study the  $\exists$ NE problem. We show that  $\exists$ NE is complete for  $\Sigma_2^P$  – the second level of the polynomial hierarchy. Namely, decision problems solvable in polynomial time by a nondeterministic Turing machine augmented by an *oracle* for an NP-complete problem. An oracle for a computational problem is a black box that is able to produce a solution for any instance of the problem in a single operation. Thus, for every problem  $P \in \Sigma_2^P$  there is a machine such that for every  $x \in P$  there is a polynomial accepting computation (with polynomial many queries to the oracle). As co-NP is the dual complexity class of NP, the dual complexity class of  $\Sigma_2^P$  is  $\Pi_2^P$ . Thus, a problem  $P$  is  $\Sigma_2^P$ -complete iff its complement  $\bar{P}$  is  $\Pi_2^P$ -complete.

**Theorem 10.** *The  $\exists$ NE problem is  $\Sigma_2^P$ -complete.*

**Proof:** For the upper bound, we describe a nondeterministic Turing machine with an oracle to SBR problem – the strict version of the BR problem, where we seek a design whose cost is strictly smaller than  $\mu$ . Given a CLG  $\mathcal{G} = \langle \mathcal{L}, \mathcal{S}_1, \dots, \mathcal{S}_k \rangle$ , we guess a profile  $P = \langle \mathcal{D}_1, \dots, \mathcal{D}_k \rangle$ , where for  $1 \leq i \leq k$ , the design  $\mathcal{D}_i$  has at most  $|\mathcal{C}_0 \times \mathcal{S}_i|$  states, where  $\mathcal{S}_i$  are the states of  $\mathcal{S}_i$ . First, we check whether  $P$  is a profile of correct designs. That is, for  $i \in [k]$ , we check whether  $\mathcal{D}_i$  is a correct design with respect to  $\mathcal{S}_i$ , which can be done in polynomial time. If there is an incorrect design we terminate and reject. Next, we check whether  $P$  is a NE profile by checking if there is a player that has a beneficial move from  $P$ . That is, for  $i \in [k]$ , we feed the oracle the input  $\langle \mathcal{G}, P, i, \text{cost}_i(P) \rangle$ . If the oracle answers YES, then Player  $i$  can benefit from deviating and  $P$  is not a NE in which case we reject. On the other hand, if for every  $i \in [k]$  the oracle answers NO, then  $P$  is a NE in which case we accept. Clearly the machine recognizes  $\exists$ NE. Note that if  $\mathcal{G} \in \exists$ NE, one of the profiles  $P$  we guess is a NE, and the computation in which we guess  $P$  is a polynomial accepting computation.

For the lower bound, we show a reduction from the complement of the  $\Pi_2^P$ -complete problem *min-max vertex cover* [20] (MMVC, for short). The input to the MMVC problem is  $\langle G, I, J, N, \mu \rangle$ , where  $G = \langle V, E \rangle$  is an undirected graph,  $I$  and  $J$  are sets of indices,  $N : V \rightarrow \{V_{i,j} \subseteq V : i \in I \text{ and } j \in J\}$  partitions the vertices, and  $\mu \in \mathbb{N}$  is a value. Note that since  $G$  is undirected, its edges are sets with two vertices. We



refer to the sets in the partition  $\{V_{i,j}\}_{i \in I, j \in J}$  as *neighborhoods* and for  $v \in V$  we refer to  $N(v)$  as the neighborhood of  $v$ . Note that there is a coarser partition of  $V$ , namely  $\{V_i\}_{i \in I}$ , where  $V_i = \bigcup_{j \in J} V_{i,j}$ . We refer to the elements in this partition as *districts* and, for  $v \in V$ , use  $D(v)$  to denote the district  $v$  belongs to. For a function  $t : I \rightarrow J$  we define  $V_t = \bigcup_{i \in I} V_{i,t(i)}$ . Intuitively,  $t$  is a choice of neighborhood in each district. Let  $G_t = \langle V_t, E_t \rangle$  be the induced subgraph of  $G$  on the vertex set  $V_t$ . Formally, for  $e \in E$ , we have  $e \in E_t$  iff  $e \subseteq V_t$ . For a graph  $G$ , we say that  $V' \subseteq V$  is a vertex cover of  $G$  if for every  $e = \{u, v\} \in E$  we have  $V' \cap \{u, v\} \neq \emptyset$ . An input  $\langle G, I, J, N, \mu \rangle$  is in MMVC iff for any choice of neighborhoods in the districts given by a function  $t$ , the smallest vertex cover of the resulting graph is at most  $\mu$ . Formally,  $\max_{t \in J^I} \min\{|V'| : V' \subseteq V_t \text{ is a vertex cover of } G_t\} \leq \mu$ . We assume without loss of generality that  $\mu \leq |V|$ .

Consider an input  $\langle G, I, J, N, \mu \rangle$  to MMVC. We construct a CLG  $\mathcal{G}$  with library  $\mathcal{L}$  and specifications  $\mathcal{S}_0, \mathcal{S}_1, \dots, \mathcal{S}_{|I|}$  such that  $\mathcal{G}$  has a NE iff  $\langle G, I, J, N, \mu \rangle \notin \text{MMVC}$ . The alphabet  $\Sigma$  consists of letters  $i, \tilde{i} \in I, v, \tilde{v} \in V$ , and  $e \in E$ . Let  $E = \{e_1, \dots, e_m\}$ . The specification of Player 0 consists of words of length  $3|E|$  of the form  $v_1 \tilde{v}_1 e_1 \dots v_m \tilde{v}_m e_m$ , for some  $v_1, \dots, v_m \in V$  (allowing duplicates). For  $i \in I$ , the specifications of Player  $i$  consist of the single word  $i \cdot (\tilde{i})^\ell$ , where  $\ell$  is a polynomial in  $|V|$ , which we define in the following.

We describe the components of  $\mathcal{L}$ . When describing the components' costs we only refer to the construction cost as their quality costs are 0. The simplest components are  $\mathcal{B}_{acc}$  and  $\mathcal{B}_{rej}$ . They consist of a single initial state that is accepting in  $\mathcal{B}_{acc}$  and rejecting in  $\mathcal{B}_{rej}$ . The cost of both these components is 0. The next component is  $\mathcal{B}_0$ , which is exactly  $\mathcal{S}_0$ , and its cost is  $\mu + 1$ . The rest of the components have no accepting states. We describe these components by the words that they can process. For each word there is a unique disjoint path from the initial to a separate exit state. For every neighborhood  $V_{i,j}$  there is a component  $\mathcal{B}_{i,j} \in \mathcal{L}$  that costs  $(3|E| + 2)(\mu + 1)$ . We refer to these components as *neighborhood components*. The single-lettered words it can process are  $i, \tilde{i}$ , and  $v$  for  $v \in V_{i,j}$ . Also, it can process the words  $v\tilde{v}e$  for  $v \in V \setminus V_{i,j}$  and  $e \in E$ , and  $\tilde{v}e$  for  $v \in V_{i,j}$  and  $e \in E$  such that  $e \cap (V_i \setminus V_{i,j}) \neq \emptyset$ . For every  $v \in V$  there is a component  $\mathcal{B}_v$  that costs 1. We refer to these components as *vertex components*. The words it can process are  $\tilde{u}e$  for  $u \in N(v)$  and  $e \in E$  such that  $v \in e$ . For  $i \in I$  there is a component  $\mathcal{B}_i$  that can process the word  $\tilde{i}$  and costs a very small value  $\xi > 0$ . The construction is clearly polynomial in the input.

We claim that  $\langle G, I, J, N, \mu \rangle \notin \text{MMVC}$  iff  $\mathcal{G} = \langle \mathcal{L}, \mathcal{S}_0, \mathcal{S}_1, \dots, \mathcal{S}_{|I|} \rangle \in \exists \text{NE}$ . For the first direction, assume  $\langle G, I, J, N, \mu \rangle \notin \text{MMVC}$ . Thus, there is a function  $t : I \rightarrow J$  such that every vertex cover of  $V_t$  has more than  $\mu$  vertices. We claim that the CLG  $\mathcal{G}$  has a NE. Consider the profile  $P$  in which Player 0 uses only the component  $\mathcal{B}_0$  and, for  $i \in I$ , Player  $i$  uses the design  $\mathcal{D}_{i,t(i)}$  that only uses the neighborhood component  $\mathcal{B}_{i,j}$  and the components  $\mathcal{B}_{acc}$  and  $\mathcal{B}_{rej}$ , which we describe in the following. For  $i \in I$  and  $j \in J$ , we describe the design  $\mathcal{D}_{i,j}$ . In  $\mathcal{D}_{i,j}$ , the component  $\mathcal{B}_{i,j}$  gains initial control. If the component relinquishes control after reading the single-lettered word  $i$ , it regains control. If it relinquishes control after reading any other word, the design assigns control to  $\mathcal{B}_{rej}$ . Similarly, for  $\ell$  times, control is given to  $\mathcal{B}_{i,j}$  assuming it reads the word  $\tilde{i}$ ,

and otherwise control is given to  $\mathcal{B}_{rej}$ . After  $\mathcal{B}_{i,j}$  gains control  $\ell + 1$  times, control is assigned to  $\mathcal{B}_{acc}$ . Clearly,  $L(\mathcal{D}_{i,j}) = L(\mathcal{S}_i)$ , thus it is a correct design for Player  $i$ .

Assume towards contradiction that  $P$  is not a NE. Thus, there is a player that benefits from deviating. We start by showing that, for  $i \in I$ , Player  $i$  cannot benefit from deviating. Consider a design  $\mathcal{D}_{i,j}$  for some  $j \neq t(i)$ . Since all the other players do not use any of the neighborhood component  $\mathcal{B}_{i,j}$  in  $P$ , we have  $cost_i(P) = cost_i(P[i \leftarrow \mathcal{D}_{i,j}])$ , thus Player  $i$  does not deviate to  $\mathcal{D}_{i,j}$ . Similarly, in every other correct design Player  $i$  uses at least two components with positive cost, which he covers alone, thus these designs cost more than  $cost_i(P)$  and he clearly does not deviate to them.

We continue to show that Player 0 cannot benefit from a deviation. Assume towards contradiction that there is a correct design  $\mathcal{D}$  such that  $cost_0(P) > cost_0(P[0 \leftarrow \mathcal{D}])$ . Consider the set  $V' \subseteq V$  of vertices that correspond to vertex components that are used in  $\mathcal{D}$ , thus  $v \in V'$  iff  $\mathcal{B}_v \in used(\mathcal{D})$ . Since  $cost_0(P) = \mu + 1$  and Player 0 does not share any of the vertex components with the players  $1, \dots, |I|$ ,  $\mathcal{D}$  uses at most  $\mu$  of them, thus  $|V'| \leq \mu$ .

We claim that there is  $V'' \subseteq V'$  such that  $V'' \subseteq V_t$  and that it is a vertex cover for  $G_t$ . Recall that the cost of using a neighborhood component without sharing is more than  $\mu + 1$ . Since we assume Player 0 benefits from deviating to  $\mathcal{D}$ , he must share all his uses of these components. Since the neighborhood components that are used by the players  $1, \dots, |I|$  are exactly these that are dictated by  $t$ , every neighborhood component  $\mathcal{B}_{i,j} \in used(\mathcal{D})$  has  $j = t(i)$ . We claim that for every  $e \in E_t$  there is a vertex  $v \in V' \cap V_t$  such that  $v \in e$ . Let  $E_t = \{e_{i_1}, \dots, e_{i_l}\}$  and  $j \in [l]$ . Consider the word  $w = v\tilde{v}e_1v\tilde{v}e_2 \dots v\tilde{v}e_{i_j-1}$  for some  $v \in V$ . Since  $w$  is a prefix of a word in  $L(\mathcal{S}_0)$ , the run of  $\mathcal{A}_{\mathcal{L},\mathcal{D}}$  on  $w$  does not get stuck. Since the last letter in  $w$  is in  $E$  the component that gains control after reading  $w$  is a neighborhood component, thus it is  $\mathcal{B}_{i,t(i)}$  for some  $i \in I$ . Consider the word  $w' = w \cdot v\tilde{v}e_{i_j}$  for  $u \in V_{i,t(i)}$ . Again, since  $w'$  is a prefix of a word in  $L(\mathcal{S}_0)$  the run of  $\mathcal{A}_{\mathcal{L},\mathcal{D}}$  on  $w'$  does not get stuck. Recall that  $\mathcal{B}_{i,t(i)}$  relinquishes control after reading  $u$ . Further recall that the components that can process words that start with  $\tilde{u}$  are  $\mathcal{B}_{i,t(i)}$  and every vertex component  $\mathcal{B}_v$  with  $v \in V_{i,t(i)}$ . Note that  $\mathcal{B}_{i,t(i)}$  cannot process the word  $\tilde{u}e_{i_j}$ . Indeed, since  $e_{i_j} \in E_t$ , both its ends are in  $V_t$ , thus there is no  $u \in V_i \setminus V_{i,t(i)}$  with  $u \in e_{i_j}$ . Thus, the components that gains control after  $\mathcal{B}_{i,t(i)}$  is a vertex component  $C_v$  with  $v \in V_{i,t(i)}$ . Since it must be able to process the word  $\tilde{u}e_{i_j}$ , we have  $v \in e_{i_j}$ , and we are done.

For the second direction, assume  $\langle G, I, J, N, \mu \rangle \in \text{MMVC}$ . Assume towards contradiction that there is a NE profile  $P$  in  $\mathcal{G}$ . We distinguish between two cases. In the first case, Player 0 does not use  $\mathcal{B}_0$ . We show that in such a case there is a player in  $I$  that benefits from deviating from  $P$ . Clearly, Player 0 must use at least one neighborhood component  $\mathcal{B}_{i,j}$ . Recall that neighborhood components cost  $(3|E|+2)(\mu+1)$ ,  $C_0$  costs  $\mu+1$ , and  $C_i$  costs  $\xi > 0$ . Since  $P$  is a NE,  $cost_0(P) \leq \mu + 1$ . Thus, Player  $i$ , the only player that can use  $\mathcal{B}_{i,j}$ , uses  $\mathcal{B}_{i,j}$  at least  $3|E|+1$  times. It is not hard to see that since  $P$  is a NE profile, Player 0 uses  $\mathcal{B}_{i,j}$  at most  $3|E|$  times, thus  $cost_i(P) \geq \frac{3|E|+1}{3|E|+3|E|+1} \cdot cost(\mathcal{B}_{i,j})$ . Consider the design  $\mathcal{D}$  that assigns initial control to  $\mathcal{B}_{i,j}$ . If it relinquishes control after reading anything different from  $i$ ,  $\mathcal{D}$  assigns control to  $\mathcal{B}_{rej}$ . Otherwise, the component  $\mathcal{B}_i$ , which can process only the word  $\tilde{i}$ , gains control  $\ell$  consecutive times after which  $\mathcal{B}_{acc}$  gains control. It is not hard to see that  $L(\mathcal{D}) = L(\mathcal{S}_i)$ . Since  $nused(\mathcal{D}, \mathcal{B}_{i,j}) = 1$

and  $nused(\mathcal{D}, \mathcal{B}_i) = \ell$ , we have  $cost_i(P[i \leftarrow \mathcal{D}]) \leq \frac{1}{2}cost(\mathcal{B}_{i,j}) + \xi$ . For sufficiently small  $\xi$ , we have  $cost_i(P) > cost_i(P[i \leftarrow \mathcal{D}])$ , and we are done.

In the second case, the design Player 0 chooses in  $P$  is the design that uses only the component  $\mathcal{B}_0$ . Thus,  $cost_0(P) = \mu + 1$ . Recall that for every  $i \in I$  and  $j \in J$ , the design  $\mathcal{D}_{i,j}$  is a correct design for Player  $i$  that uses only the neighborhood component  $\mathcal{B}_{i,j}$ . It is not hard to see that since  $P$  is a NE, every Player  $i$  chooses some design  $\mathcal{D}_{i,j}$ . Let  $t : I \rightarrow J$  be the function that corresponds to these players choices. Since  $\langle G, I, J, N, \mu \rangle \in \text{MMVC}$ , there is a vertex cover  $V' \subseteq V_t$  of  $G_t$  such that  $|V'| \leq \mu$ .

We construct a design  $\mathcal{D}$  for Player 0 such that  $cost_0(P[0 \leftarrow \mathcal{D}]) < \mu + 1$ , thus we reach a contradiction to the fact that  $P$  is a NE. Recall that Player 0's specification is the set of words of length  $3|E|$  of the form  $v_1 \tilde{v}_1 e_1 \dots v_m \tilde{v}_m e_m$ , where  $E = \{e_1, \dots, e_m\}$  and  $v_1, \dots, v_m \in V$  (allowing duplicates). The definition of  $\mathcal{D}$  is inductive. Let  $1 \leq l \leq |E|$ . Consider a word  $w \in \Sigma^*$  of length  $3(l-1)$  that can be extended to a word in  $L(\mathcal{S}_0)$ . That is, there is a word  $x \in \Sigma^*$  such that  $w \cdot x \in L(\mathcal{S}_0)$ . For  $v \in V$ , let  $w_v = w \cdot v \tilde{v} e_{l+1}$ . Note that  $w_v$  can also be extended to a word in  $L(\mathcal{S}_0)$  (possibly with  $\epsilon$ ). Assuming the run of  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  on  $w$  does not get stuck and control is relinquished from some component after reading  $w$ , we describe how  $\mathcal{D}$  assigns control next such that the run of  $\mathcal{A}_{\mathcal{L}, \mathcal{D}}$  on every  $w_v$  does not get stuck and control is relinquished after reading  $w_v$ .

We distinguish between two cases. In the first case,  $e_l \notin E_t$ . Thus, there is a vertex  $v \in e_l \setminus V_t$ . Let  $V_i$  be  $v$ 's district, thus  $V_i = D(v)$ . Let  $V_{i,j} \subseteq V_i$  be the neighborhood in  $V_i$  that selected by  $t$ , thus  $j = t(i)$ . Note that this since  $v \notin V_t$ , it does not belong to  $V_{i,j}$ . The component to which  $\mathcal{D}$  assigns control after reading  $w$  is the neighborhood component  $\mathcal{B}_{i,j}$ . Consider a vertex  $u \in V$ . If  $u \notin V_{i,j}$ , then when reading  $u \tilde{u} e_l$  the run in  $\mathcal{B}_{i,j}$  does not get stuck and control is relinquished at its end. If  $u \in V_{i,j}$ , then  $\mathcal{B}_{i,j}$  relinquishes control after reading  $u$ . In such a case we define  $\mathcal{D}$  to reassign control to  $\mathcal{B}_{i,j}$ . Note that the run of  $\mathcal{B}_{i,j}$  on  $\tilde{u} e_l$  does not get stuck. Indeed, the vertex  $v \in e_l$  is a witness to the fact that  $(V_i \setminus V_{i,j}) \cap e_l \neq \emptyset$ . Clearly, control is relinquished after  $\mathcal{B}_{i,j}$  reads  $\tilde{u} e_l$ . If in one of the times  $\mathcal{B}_{i,j}$  gains control it relinquishes it after reading any other word  $x \in \Sigma^*$ , then there is no  $y \in \Sigma^*$  such that  $w \cdot x \cdot y \in L(\mathcal{S}_0)$ , and we assign control to  $\mathcal{B}_{rej}$ .

In the second case,  $e_l \in E_t$ . Thus, there is a vertex  $v$  in the vertex cover  $V'$  such that  $v \in e_l$ . Let  $N(v) = V_{i,j}$ . Note that since  $v \in V_t$ , we have  $t(i) = j$ . The component to which  $\mathcal{D}$  assigns control after reading  $w$  is the neighborhood component  $\mathcal{B}_{i,j}$ . Consider a vertex  $u \in V$ . The case where  $u \notin V_{i,j}$  is similar to the previous case. For  $u \in V_{i,j}$ , when  $\mathcal{B}_{i,j}$  reads  $u$ , it relinquishes control. In such a case,  $\mathcal{D}$  assigns control to  $\mathcal{B}_v$ . Since  $v \in V_{i,j}$  and  $v \in e_l$ , the component  $\mathcal{B}_v$  can process the word  $\tilde{u} e_l$ , and it relinquishes control after its end. Similarly to the above, if  $\mathcal{B}_{i,j}$  or  $\mathcal{B}_v$  relinquish control after reading any other word  $x \in \Sigma^*$ , then we assign control to  $\mathcal{B}_{rej}$ . Finally, if  $l = |E|$ , we assign control  $\mathcal{B}_{acc}$ .

Correctness of the design  $\mathcal{D}$  is immediate from the construction. We claim that  $cost_0(P[0 \leftarrow \mathcal{D}]) < \mu + 1$ . Note that the  $\mathcal{D}$  uses two types of components. The first type is neighborhood components. Consider  $\mathcal{B}_{i,j} \in used(\mathcal{D})$ . Note that we constructed  $\mathcal{D}$  so that  $j = t(i)$ . Thus, Player 0 shares the cost of  $\mathcal{B}_{i,j}$  with Player  $i$ . Recall that Player  $i$  chooses the design  $\mathcal{D}_{i,j}$  that uses  $\mathcal{B}_{i,j}$   $\ell$  times. We define  $\ell$  to be sufficiently large so that

the proportion of the cost that Player 0 pays is less than  $\frac{1}{2|E|}$ . Thus, the total cost Player 0 endures for neighborhood components is less than 1. The second type of components  $\mathcal{D}$  uses is vertex components. Since  $\mathcal{D}$  uses vertex components that correspond to vertices in the vertex cover  $V'$ , the number of such component that  $\mathcal{D}$  uses is at most  $\mu$ , which is the total cost for these components. Thus,  $cost_0(P[0 \leftarrow \mathcal{D}]) < \mu + 1$ , and we are done. ■

## 6 Discussion

Traditional synthesis algorithms assumed that the system is constructed from scratch. Previous work adjusted synthesis algorithms to a reality in which systems are constructed from component libraries. We adjust the algorithms further, formalize the notions of quality and cost and seek systems of high quality and low cost. We argue that one should distinguish between quality considerations, which are independent of uses of the library by other designs, and pricing considerations, which depend on uses of the library by other designs.

Once we add multiple library users to the story, synthesis is modeled by a resource-allocation game and involves ideas and techniques from algorithmic game theory. In particular, different models for sharing the price of components can be taken. Recall that in our model, users share the price of a component, with the share being proportional to the number of uses. In some settings, a *uniform sharing rule* may fit better, which also makes the game more stable. In other settings, a more appropriate sharing rule would be the one used in *congestion games* – the more a component is used, the higher is its price, reflecting, for example, a higher load. Moreover, synthesis of different specifications in different times gives rise to *dynamic allocation* of components, and synthesis of collections of specifications by different users gives rise to *coalitions* in the games. These notions are well studied in algorithmic game theory and enable an even better modeling of the rich settings in which traditional synthesis is applied.

## References

1. M. Abadi, L. Lamport, and P. Wolper. Realizable and unrealizable concurrent program specifications. In *Proc. 25th ICALP*, LNCS 372, pages 1–17. Springer, 1989.
2. S. Almagor, U. Boker, and O. Kupferman. Formalizing and reasoning about quality. In *Proc. 40th ICALP*, LNCS 7966, pages 15 – 27. Springer, 2013.
3. G. Alonso, F. Casati, H.A. Kuno, and V. Machiraju. Web Services - Concepts, Architectures and Applications. *Data-Centric Systems and Applications*. Springer, 2004.
4. R. Alur, K. Etessami, and P. Madhusudan. A temporal logic of nested calls and returns. In *Proc. 10th TACAS*, LNCS 2725, pages 67–79. Springer, 2004.
5. G. Avni and O. Kupferman. When does abstraction help? *IPL*, 113:901–905, 2013.
6. G. Avni, O. Kupferman, and T. Tamir. Congestion and cost-sharing games with multisets of resources. *Submitted*, 2014.
7. G. Avni, O. Kupferman, and T. Tamir. Network-formation games with regular objectives. In *Proc. 17th FoSSaCS*, LNCS 8412, pages 119–133. Springer, 2014.
8. D. Berwanger and L. Doyen. On the power of imperfect information. In *Proc. 28th TST&TCS*, pages 73–82, 2008.
9. A. Bohy, V. Bruyère, E. Filiot, and J-F. Raskin. Synthesis from LTL specifications with mean-payoff objectives. In *Proc. 19th TACAS*, LNCS 7795, pages 169–184. Springer, 2013.

10. U. Boker, K. Chatterjee, T.A. Henzinger, and O. Kupferman. Temporal specifications with accumulative values. In *Proc. 26th LICS*, pages 43–52, 2011.
11. L. de Alfaro, M. Faella, T.A. Henzinger, R. Majumdar, and M. Stoelinga. Model checking discounted temporal properties. *TCS*, 345(1):139–170, 2005.
12. L. de Alfaro and T.A. Henzinger. Interface theories for component-based design. In *Proc. 1st EMSOFT*, LNCS 2211, pages 148–165. Springer, 2001.
13. L. Doyen, T. A. Henzinger, B. Jobstmann, and T. Petrov. Interface theories with component reuse. In *Proc. 8th EMSOFT*, pages 79–88, 2008.
14. J. Elgaard, N. Klarlund, and A. Möller. Mona 1.x: new techniques for WS1S and WS2S. In *Proc. 10th CAV*, LNCS 1427, pages 516–520. Springer, 1998.
15. A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. In *Proc. 36th STOC*, pages 604–612, 2004.
16. M. Faella, A. Legay, and M. Stoelinga. Model checking quantitative linear time logic. *ENTCS*, 220(3):61–77, 2008.
17. E. Filiot, N. Jin, and J.-F. Raskin. Antichains and compositional algorithms for LTL synthesis. *FMSD*, 39(3):261–296, 2011.
18. G. Göbller and J. Sifakis. Composition for component-based modeling. *Sci. Comput. Program.*, 55(1-3):161–183, 2005.
19. M. Jurdzinski. Small progress measures for solving parity games. In *Proc. 17th STACS*, LNCS 1770, pages 290–301. Springer, 2000.
20. K.-I. Ko and C.-L. Lin. On the complexity of min-max optimization problems and their approximation. In *Minimax and Applications*, volume 4 of *Nonconvex Optimization and Its Applications*, pages 219–239. Springer, 1995.
21. O. Kupferman, N. Piterman, and M.Y. Vardi. Safraless compositional synthesis. In *Proc. 18th CAV*, LNCS 4144, pages 31–44. Springer, 2006.
22. O. Kupferman and M.Y. Vardi. Safraless decision procedures. In *Proc. 46th FOCS*, pages 531–540, 2005.
23. Y. Lustig and M.Y. Vardi. Synthesis from component libraries. *STTT* 15:603–618, 2013.
24. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th POPL*, pages 179–190, 1989.
25. J.-F. Raskin, K. Chatterjee, L. Doyen, and T. Henzinger. Algorithms for  $\omega$ -regular games with imperfect information. *LMCS*, 3(3), 2007.
26. T. Roughgarden and E. Tardos. How bad is selfish routing? *JACM*, 49(2):236–259, 2002.
27. S. Safra. On the complexity of  $\omega$ -automata. In *Proc. 29th FOCS*, pages 319–327, 1988.
28. A.P. Sistla and E.M. Clarke. The complexity of propositional linear temporal logic. *JACM*, 32:733–749, 1985.