
SEQUENTIAL RELATIONAL DECOMPOSITION

DROR FRIED, AXEL LEGAY, JOËL OUAKNINE, AND MOSHE Y. VARDI

Department of Computer Science, The Open University of Israel, Israel
e-mail address: dfried@openu.ac.il

Department of Computer Science, Universit Catholique de Louvain, Belgium
e-mail address: axel.legay@uclouvain.be

Max Planck Institute for Software Systems, Saarland Informatics Campus, Germany

Department of Computer Science, Oxford University, UK
e-mail address: joel@mpi-sws.org

Department of Computer Science, Rice University, USA
e-mail address: vardi@cs.rice.edu

ABSTRACT. The concept of *decomposition* in computer science and engineering is considered a fundamental component of *computational thinking* and is prevalent in design of algorithms, software construction, hardware design, and more. We propose a simple and natural formalization of *sequential decomposition*, in which a task is decomposed into two sequential sub-tasks, with the first sub-task to be executed out before the second sub-task is executed. These tasks are specified by means of input/output relations. We define and study *decomposition problems*, which is to decide whether a given specification can be sequentially decomposed. Our main result is that decomposition itself is a difficult computational problem. More specifically, we study decomposition problems in three settings: where the input task is specified explicitly, by means of Boolean circuits, and by means of automatic relations. We show that in the first setting decomposition is NP-complete, in the second setting it is NEXPTIME-complete, and in the third setting there is evidence to suggest that it is undecidable. Our results indicate that the intuitive idea of decomposition as a system-design approach requires further investigation. In particular, we show that adding human to the loop by asking for a decomposition hint lowers the complexity of decomposition problems considerably.

1. INTRODUCTION

Over the past decade, it became apparent that the conceptual way of analyzing problems through computer-science techniques can be considered as a general approach to analysis, design, and problem solving, known as *computational thinking* [44, 45]. A key element of this approach is the taming of complexity by decomposing a complex problem into simpler

Key words and phrases: Decomposition, Composition, Automatic Relations, Positivity, Synthesis.

An extended abstract of this article appeared in the Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018.

problems. Quoting Wing [44]: “Computational thinking is using abstraction and decomposition when attacking a large complex task or designing a large complex system.” While abstraction helps to tame complexity by simplifying away irrelevant details of a complex problem, decomposition helps complexity by breaking down complex problems into simpler ones. In fact, decomposition is a generic project-management technique: project managers quite often face challenges that require decomposition – a large project that is divided among team members to make the problem less daunting and easier to solve as a set of smaller tasks, where team members work on tasks that are in their specific fields of expertise. As computer scientists and engineers, the concept of decomposition is prevalent in the design of algorithms, in software construction, in hardware design, and so on. For example, a classical paper in software engineering studies criteria to be used in decomposing systems into modules [37]. Yet, in spite of the centrality of the concept of decomposition in computational thinking, it has yet to be studied formally in a general theoretical setting (see related work). Such a study is the focus of this work.

There are many different types of decomposition that can be considered. Based on her understanding of the problem, the *decomposer* has to make a decision on how to decompose a given problem, for example, by meeting certain constraints on the size of the sub-problems, or constraints on the way that solved sub-problems ought to be recomposed. A simple and natural way of decomposition is *sequential* decomposition in which a task is decomposed into two sub-tasks, where the first sub-task is to be carried out before the second sub-task can be executed. A formal model for sequential decomposition is the subject of this work. We assume that the given problem is specified by means of an input/output relation. It is widely accepted that such relations are the most general way to specify programs, whether for terminating programs [19], where input and output should be related values, or for nonterminating programs [38], where input and output are streams of values. The *decomposition problem* is to decompose a given input/output relation R , between an input domain $\mathcal{I}$ and an output domain $\mathcal{O}$, into two relations R_1 and R_2 , such that R can be reconstructed from R_1 and R_2 using relational composition (defined in Section 3). To avoid trivial solutions, where either R_1 or R_2 is the identity relation, we assume that the intermediate domain, that is, the co-domain of R_1 , which is also the domain of R_2 , is specified. Intuitively, specifying the intermediate domain amounts to constraining the manner in which the first task can “communicate” with the second task. Such a restriction can be viewed as a form of *information hiding*, which is one of the major criteria for decomposition in [37]. As we show, sequential decomposition is nontrivial only when the channel of communication between the first and second task has a small “bandwidth”, appropriately defined.

We study sequential decomposition in three settings: *explicit*, *symbolic*, and *automatic*. In the explicit setting, the input/output relation R is specified explicitly. In the symbolic setting, the domains and R are finite but too large to be specified explicitly, so R is specified symbolically as a Boolean circuit. In the automatic setting, the domains and R may be infinite, so the domains are specified by means of an alphabet, over which R is specified by means of a deterministic finite-state automaton.

Our general finding is that sequential decomposition, viewed as a computational problem, is itself a challenging problem. In the explicit setting, the decomposition problem is NP-complete. This escalates to NEXPTIME-complete for the symbolic setting. For the automatic setting the decomposition problem is still open, but we provide evidence and conjecture that it is undecidable. Specifically, we show that even a very simple variant of

the automatic setting can be viewed as equivalent to the Positivity problem, whose decidability is well known to be open [42, 34]. We do show, however, that a “strategic” variant of the automatic setting, in which the required relations are described as transducers is in EXPTIME. These findings, that decomposition is an intractable problem, can be viewed as a “No-Free-Lunch” result, as it says that decomposition, which is a tool to combat complexity, is itself challenged by computational complexity. This means that while decomposition is an essential tool, the application of decomposition is an art, rather than science, and requires human intuition.

As such, we explore decomposition with “a human in the loop”, where the role of the human is to offer a hint, suggesting one of the terms of the decomposition, and the role of the decomposition algorithm is to check if a valid decomposition can be found based on the hint. Table 1 summarizes our complexity results with and without hints. We show that in most cases decomposition with a hint is easier than decomposition with no hint.

Table 1: Computational complexity of decomposition without and with hints.

	Without hints	With hint R_1	With hint R_2
Explicit	NP-complete	PTIME	PTIME
Symbolic	NEXPTIME-complete	Π_3^P	Π_3^P
Automatic	undecidable?	EXPTIME	EXPTIME
Strategic	EXPTIME	EXPTIME	PTIME

2. RELATED WORK

In this paper we introduce a new framework to study *decomposition* of system-level specifications into component-level specifications. In contrast, most existing approaches in software engineering focus on *composition*, that is, developing systems from independently developed components, or proving system-level properties from component-level properties, for example [15]. This compositional approach is a fundamental approach in computer science to taming complexity.

We now give a few examples of the composition-based approach to system development. The *contract-theory* approach of Benveniste et al. [6] uses assumption/guarantee to describe both the implementation and the environment of a given component; that is, the implementation of the component should work with any other component whose implementation satisfies the assumption. Components are then composed in a bottom up fashion to produce a large-size system. The same approach is pursued also in the *interface theories* of de Alfaro et al. [12], the *modal automata* of Larsen [31], and the *I/O automata* of Lynch [33].

Operations such as *quotient* (adjoint of composition) or contextual equation solving [30, 13] allow one to synthesize a component that can be composed with another one in order to refine a large size specification. This gets close to the problem of decomposition with *hints*, described below, but is still focused on synthesizing an implementation. The BIP approaches of Sifakis [5] proposes to specify complex systems by aggregating smaller components using a very expressive algebra. Recently in [10], the approach has been extended to synthesize the order of communication between components, but the focus there is not on decomposing specification. Similarly, Lustig and Vardi have shown [32] how to synthesize reactive systems satisfying given linear-temporal properties by composing components from a given library of reactive components. Closer to our work is an approach

studied in model-driven software engineering. The Fragmenta tool [3] provides algebraic descriptions of models and fragments based on graph morphism. While the tool eases the decomposition task, it does not specifically study the problem of actually providing such a decomposition.

A major weakness of composition-based approaches to system development is that they consider only the problem of simplifying the implementation work, but ignore the task of decomposing the often very complex system-level technical specifications into smaller/simpler ones. For example, there is no technique to explain how smaller assumption/guarantee contracts can be obtained from larger ones. This operation has to be conducted manually by developers using their intuition and understanding of complex system-level specification. Thus, our work here on decomposition complements existing approaches on composition-based development. In addition, we believe that our work is also relevant to *architectural design*, e.g., as a complement of [7].

A classical paper of Parnas in software engineering studies criteria for to be used in decomposing systems into modules [37]. Parnas's framework, however, assumes that the starting point for decomposition consists of architectural specification of the system, while our starting point is quite more abstract, as we assume that we are provided with relational specification. Another related work is that of [40], which describe an approach for extracting sequential components from system specification. Unlike, however, our work here, which starts from a highly abstract relational specification, the approach of [40] assumes that the system's specification is provided by means of interface specification of the components. Thus, this approach is of factorization rather than decomposition.

Decomposition has been studied in the context of linear algebra. A matrix decomposition or matrix factorization is a factorization of a matrix into a product of matrices. There are many different matrix factorizations. Certain Boolean matrix-factorization problems are known to be NP-complete [22]. Our NP-completeness result for explicit relations can be viewed as a special case of Boolean matrix-factorization. In addition, our formulation for the explicit case can be viewed as a reformulation of the combinatorial definition of the nondeterministic communication complexity (see Chapters 1-2 in [28]). In that sense, this paper extends these works to more general representations of relations.

3. PRELIMINARIES

3.0.1. Relations. Let A, B, C be finite sets. For a binary relation $R \subseteq A \times B$, let $Dom(R)$, and $Img(R)$ be the domain of R , and the image (sometimes called co-domain) of R , defined as follows. $Dom(R) = \{a \in A \mid \exists b \in B \text{ s.t. } (a, b) \in R\}$, and $Img(R) = \{b \in B \mid \exists a \in A \text{ s.t. } (a, b) \in R\}$. For $a \in A$, let $Img_a(R) = \{b \in B \mid (a, b) \in R\}$. R is called a function if for every $a \in A, b, b' \in B$ we have $(a, b), (a, b') \in R \implies b = b'$. Given binary relations $R_1 \subseteq A \times B$, and $R_2 \subseteq B \times C$, the *composition* of R_1 and R_2 is a binary relation $R_1 \circ R_2 \subseteq A \times C$ where $R_1 \circ R_2 = \{(a, c) \mid \exists b \in B \text{ s.t. } (a, b) \in R_1 \text{ and } (b, c) \in R_2\}$.

3.0.2. Automata. A Nondeterministic Finite Automaton (NFA) is a tuple $\mathcal{A} = (\Sigma, Q, q, \delta, F)$, where Σ is a finite alphabet, Q is a finite state set with an initial state q , $\delta : Q \times \Sigma \rightarrow 2^Q$ is a transition function, and F is an accepting-state set. A run of $\mathcal{A}$ over a word $w = a_1 a_2 \cdots a_n$ for some n is a state sequence $r = q_0, q_1, \dots, q_n$ such that $q_{i+1} \in \delta(q_i, a_i)$ for $i \geq 0$, where $q_0 = q$ is the initial state. A run is *accepting* if its final state is accepting. A

word is *accepted* if it has an accepting run. The language of $\mathcal{A}$, $L(\mathcal{A})$, is the set of all accepted words of $\mathcal{A}$ and is called a *regular language*. A language is also regular if and only if it can be described by a *regular expression*. We define the size of the automaton $\mathcal{A}$ as $|Q| + |\Sigma| + |\delta|$ and denote this size by $|\mathcal{A}|$. For NFAs $\mathcal{A}_1 = (\Sigma_1, Q_1, q^1, \delta_1, F_1)$ and $\mathcal{A}_2 = (\Sigma_2, Q_2, q^2, \delta_2, F_2)$, we define the *product automaton* of $\mathcal{A}_1$ and $\mathcal{A}_2$ as an automaton $\mathcal{A}_1 \times \mathcal{A}_2 = (\Sigma_1 \times \Sigma_2, Q_1 \times Q_2, q^1 \times q^2, \delta, F_1 \times F_2)$ where $(p, p') \in \delta((q, q'), (l, l'))$ iff $p \in \delta_1(q, l)$ and $p' \in \delta_2(q', l')$. An NFA $\mathcal{A}$ is deterministic (called DFA) if for every state q and letter a , $|\delta(q, a)| \leq 1$. Every NFA can be determinized to a DFA that describes the same language by using the *subset construction*, possibly with an exponential blow-up [20]. It is often more convenient for users to specify regular languages by means of regular expressions, which can be converted to DFA, possibly with a doubly exponential blow-up [20]. We assume here that all regular languages are specified by means of DFAs, as we wish to study the inherent complexity of decomposition.

Finally a *transducer* (also called a Moore machine) is a DFA with no accepting states, but with additional output alphabet and an additional function from the set of states to the output alphabet. Formally a transducer is a tuple $T = (\Sigma, \Sigma', Q, \delta, q, Out)$ where Σ is a finite alphabet and Σ' is a finite output alphabet, Q a set of states with an initial state $q \in Q$, $\delta : Q \times \Sigma \rightarrow Q$ is the transition function and $Out : Q \rightarrow \Sigma'$ is the output function. Transducers describe automatic functions from input to output.

4. PROBLEM DEFINITION

The concept of decompositions that we explore here is related to systems that can be defined by their given input and produced output. We model these as an input domain $\mathcal{I}$ and an output domain $\mathcal{O}$, not necessarily finite. Our description of a system is a specification that associates inputs to outputs, and is modeled as a relation $R \subseteq \mathcal{I} \times \mathcal{O}$ [19, 38]. In addition we assume a constraint in form of a domain with a specific size that directs the decomposition to be more concise and is given as an intermediate domain $\mathcal{B}$. The objective is to decompose R into relations $R_1 \subseteq \mathcal{I} \times \mathcal{B}$, and $R_2 \subseteq \mathcal{B} \times \mathcal{O}$ such that the composition $R_1 \circ R_2$ has either of the following properties: (i) no input-output association is added or lost - this problem is called the *Total Decomposition Problem (TDP)*, or (ii) a more relaxed version, called the *Partial Decomposition Problem (PDP)* in which no input-output association is added, but we are allowed to lose some of the output as long as each input can be resolved.

To make the paper more fluent to read we use the notation TDP/PDP for statements that are valid to the TDP and the PDP variants respectively. Since the size of the intermediate domain $\mathcal{B}$ can be significantly smaller than the size of the input or output domains, the problem becomes non-trivial as some sort of compression is required in order to enable TDP/PDP. The actual problems of TDP/PDP are appropriately defined for each section as decision problems. We first define the TDP/PDP conditions as follows.

Definition 4.1. (TDP/PDP conditions) Given binary relations $R \subseteq \mathcal{I} \times \mathcal{O}$, $R_1 \subseteq \mathcal{I} \times \mathcal{B}$, and $R_2 \subseteq \mathcal{B} \times \mathcal{O}$ for some domains $\mathcal{I}, \mathcal{O}, \mathcal{B}$, we say that (R_1, R_2) meet the TDP condition if $Img(R_1) \subseteq Dom(R_2)$ and $R_1 \circ R_2 = R$. We say that (R_1, R_2) meet the PDP condition if $Img(R_1) \subseteq Dom(R_2)$, $Dom(R_1 \circ R_2) = Dom(R)$, and $R_1 \circ R_2 \subseteq R$.

The decision problem of TDP/PDP, formally defined for diverse settings, is: given a description of domains and a relevant relation R , find whether there exist (R_1, R_2) that meet the TDP/PDP condition.

Since the type of domains and relations that we explore varies between an explicit and a more implicit description of relations, the formal definitions of the problem change according to these representations, and so are the sought decomposed relations. It is important to note that the TDP/ PDP conditions are properties of the actual relations, and not of the description in which the relations are represented.

Note that PDP without the restriction of $Dom(R_1 \circ R_2) = Dom(R)$ becomes trivial, as one can take the empty sets as R_1 and R_2 .

The following technical fact can help prove TDP/PDP conditions in various settings.

Claim 4.2. Assume $Img(R_1) \subseteq Dom(R_2)$. Then $Dom(R_1 \circ R_2) = Dom(R_1)$ (and therefore $Dom(R) = Dom(R_1)$).

Proof. Let $i \in Dom(R_1 \circ R_2)$. Then there is $o \in \mathcal{O}$ such that $(i, o) \in R_1 \circ R_2$. Which means there is $b \in \mathcal{B}$ such that $(i, b) \in R_1$ and $(b, o) \in R_2$. Specifically there is $b \in \mathcal{B}$ such that $(i, b) \in R_1$, therefore $i \in Dom(R_1)$. Next let $i \in Dom(R_1)$. Then there is $b \in \mathcal{B}$ such that $(i, b) \in R_1$. Since $Img(R_1) \subseteq Dom(R_2)$, we have that $b \in Dom(R_2)$ therefore there is $o \in \mathcal{O}$ such that $(b, o) \in R_2$. Therefore since $(i, b) \in R_1$ and $(b, o) \in R_2$ then $(i, o) \in R_1 \circ R_2$ which makes $i \in Dom(R_1 \circ R_2)$. $\square$

Finally see that the decomposed relations that meet the TDP conditions, also meet the PDP conditions on the same input, therefore a positive answer to TDP implies a positive answer to PDP.

5. DECOMPOSITION IS HARD

Decomposition has been advocated as the first step in the design of complex systems, with the intuition that it is easier to design components separately, rather than design a complex monolithic system. We show in this section several settings in which finding a sequential decomposition is computationally hard. This means that while sequential decomposition could be used to simplify the complexity of the initial specification, such decomposition itself is intractable, thus can be viewed as a "No-Free-Lunch".

5.1. Explicit relations. The simplest case of decomposition is when the domains $\mathcal{I}, \mathcal{O}$ and $\mathcal{B}$ are finite and given explicitly as a part of the input, and the relation R is given explicitly as a table in $\mathcal{I} \times \mathcal{O}$.

Problem 5.1. (TDP/PDP on explicit relations) We are given a tuple $I = (\mathcal{I}, \mathcal{O}, \mathcal{B}, R)$ where $\mathcal{I}, \mathcal{O}, \mathcal{B}$ are finite domains and $R \subseteq \mathcal{I} \times \mathcal{O}$. The problem is whether there exist relations $R_1 \subseteq \mathcal{I} \times \mathcal{B}$, and $R_2 \subseteq \mathcal{B} \times \mathcal{O}$ such that (R_1, R_2) meet the TDP/PDP conditions.

Claim 5.2. If $|\mathcal{B}| \geq |\mathcal{I}|$ or $|\mathcal{B}| \geq |\mathcal{O}|$ then TDP has a positive solution (and therefore PDP as well) and the relations that solve TDP can be found in a linear time to the size of the input.

Proof. Suppose w.l.o.g. that $|\mathcal{B}| \geq |\mathcal{O}|$. Then we can define $R_2 = \{(g(o), o) | o \in \mathcal{O}\}$ where $g : \mathcal{O} \rightarrow \mathcal{B}$ is any injection. Then for TDP, $R_1 = \{(i, b) | (i, o) \in R \text{ and } b = g(o)\}$ is a relation that satisfies $R_1 \circ R_2 = R$. $\square$

Therefore TDP/PDP become non-trivial when $|\mathcal{B}|$ is strictly smaller than $|\mathcal{I}|$ and $|\mathcal{O}|$.

Example 5.3. Let $\mathcal{I} = \{i_1, i_2\}$, $\mathcal{B} = \{b\}$, and $\mathcal{O} = \{o_1, o_2\}$. Let $R = \{(i_1, o_1), (i_2, o_2)\}$. Then the answer to TDP is negative as every non-empty composition of relations $R_1 \circ R_2$ with $\text{Dom}(R_1 \circ R_2) = \text{Dom}(R)$, must also include (i_1, o_2) or (i_2, o_1) .

We next show that even for explicit setting, TDP/PDP are computationally hard, that is NP-complete. On a positive note, being in NP, solutions for TDP/PDP can be sought by various techniques such as reduction to SAT, then using SAT solvers.

Theorem 5.4. *TDP/PDP on explicit relations are NP-complete.*

Proof. TDP on explicit relations is NP-complete. To see that TDP is in NP, guess R_1, R_2 and verify TDP conditions. We show hardness by reduction from the NP-complete problem: Covering by Complete Bipartite Subgraphs (CCBS) (Problem GT18 at [18]). In the CCBS we are given a bipartite graph G and $k > 0$, and the problem is whether G can be covered by k complete bipartite subgraphs.

Given a CCBS instance $G = (V_1, V_2, E)$ and $k > 0$ we define a TDP instance with $\mathcal{I} = V_1$ and $\mathcal{O} = V_2$ (w.l.o.g. every vertex in V_1 has an outgoing edge, and every vertex in V_2 has an incoming edge). Set $R = E$ and add the $\mathcal{B}$ elements $\mathcal{B} = \{u_1, \dots, u_k\}$. Suppose there is a solution (R_1, R_2) to TDP. For every $i \leq k$ let $V_1^i = \{v \in V_1 \mid (v, u_i) \in R_1\}$, and let $V_2^i = \{v \in V_2 \mid (u_i, v) \in R_2\}$. From TDP requirement it follows that $(v, v') \in R$ for every $v \in V_1^i, v' \in V_2^i$. Furthermore, for every $(v, v') \in R$ there is $i \leq k$ such that $v \in V_1^i, v' \in V_2^i$. Since $E = R$, then $\{(V_1^i, V_2^i) \mid i \leq k\}$ is a collection of complete bipartite subgraphs that covers G .

Next, assume there is a solution $\{(V_1^i, V_2^i) \mid i \leq k\}$ to the CCBS. Then define $R_1 = \bigcup_{i \leq k} \{(v, u_i) \mid v \in V_1^i\}$, and $R_2 = \bigcup_{i \leq k} \{(u_i, v) \mid v \in V_2^i\}$. Then $R_1 \subseteq \mathcal{I} \times \mathcal{B}$, $R_2 \subseteq \mathcal{B} \times \mathcal{O}$ and $\text{Img}(R_1) \subseteq \text{Dom}(R_2)$. To see that $R_1 \circ R_2 = R$, let $(v, v') \in R_1 \circ R_2$. Then there is u_i such that $(v, u_i) \in R_1, (u_i, v') \in R_2$, so $(v, v') \in (V_1^i, V_2^i)$, therefore $(v, v') \in R$. On the other hand let $(v, v') \in R$. Since the solution to CCBS covers G , there is $i \leq k$ such that $(v, v') \in (V_1^i, V_2^i)$, so $(v, u_i) \in R_1$ and $(u_i, v') \in R_2$, hence $(v, v') \in R_1 \circ R_2$.

PDP on explicit relations is NP-complete. Membership in NP is easily shown. We show hardness by reduction from Set Cover (Problem SP5 in [18]). In Set Cover we are given a set of elements $S = \{a_1, \dots, a_n\}$, a set of subsets of S , $C = \{c_1, \dots, c_m\}$ and $k \geq 0$. The problem is whether there are k sets from C that their union covers S . We assume w.l.o.g. that every member of S belongs to at least one member of C . Given an instance $I = (S, C, k)$ of Set Cover, define a PDP instance I' where $\mathcal{I} = S$, $\mathcal{O} = C$, $R = \{(a, c) \mid a \in c\}$ and add the intermediate domain elements to be $\mathcal{B} = \{b_1, \dots, b_k\}$. Assume that there is a set cover $\{c_{i_1}, \dots, c_{i_k}\} \subseteq C$ for S . Then construct R_1, R_2 as follows. For every element a and $j \leq k$, set $(a, b_j) \in R_1$ iff c_{i_j} is the first-most set for which $a \in c_{i_j}$. In addition set $(b_j, c_{i_j}) \in R_2$. Then $R_1 \subseteq \mathcal{I} \times \mathcal{B}$, $R_2 \subseteq \mathcal{B} \times \mathcal{O}$ and $\text{Img}(R_1) \subseteq \text{Dom}(R_2)$. Since the set cover for S covers all elements of S , we have that $\text{Dom}(R_1 \circ R_2) = \text{Dom}(R)$. To see that $R_1 \circ R_2 \subseteq R$, let $(a, c_{i_j}) \in R_1 \circ R_2$ for some $j \leq k$. Then $(a, b_j) \in R_1$ and $(b_j, c_{i_j}) \in R_2$. By definition of R_1 , $a \in c_{i_j}$, therefore we have that $(a, c_{i_j}) \in R$.

Next, assume I' has a positive PDP solution with relations (R_1, R_2) . Define a set cover as follows. For every $b_j \in \text{Img}(R_1)$ (and therefore $b_j \in \text{Dom}(R_2)$), set a single element $c_{i_j} \in \text{Im}g_{b_j}(R_2)$ and set $c_{i_j} \in C'$. Then $|C'| \leq k$. To see that C' is a set cover for S , let $a \in \mathcal{I}$. Then since $a \in \text{Dom}(R)$ and $\text{Dom}(R_1 \circ R_2) = \text{Dom}(R)$, there is $c \in \mathcal{O}$ such that $(a, c) \in R_1 \circ R_2$. Therefore there is $b_j \in \mathcal{B}$ such that $(a, b_j) \in R_1$ and by definition

$(b_j, c_{i_j}) \in R_2$, which means $(a, c_{i_j}) \in R_1 \circ R_2$ as well. Since $R_1 \circ R_2 \subseteq R$ then $(a, c_{i_j}) \in R$ which means $a \in c_{i_j}$. As $c_{i_j} \in C'$ our proof is complete. $\square$

5.2. Symbolic relations. In Section 5.1, the input relation is described explicitly. In many cases, however, although finite, the relation is too large to be described explicitly, and it makes more sense to describe it symbolically. Specifically, the domains are given as the set of all truth assignments over sets of Boolean variables, and the relation is described symbolically. Such representations have been studied in the literature, where they are often referred to as *succinct* representations, since they allow for a polynomial-size description of exponential-size domains and relations. In this section we explore a standard encoding in which the relation is described as a Boolean circuit, as in [11, 4]. Other symbolic encodings studied in the literature are Boolean formulas [43], and BDDs [14].

Let D, D' be finite domains of size $|D| = 2^n$, and $|D'| = 2^{n'}$ for some n, n' . A (succinct) *circuit* description of a relation $R \subseteq D \times D'$ is a circuit C_R of size polynomial in $\max(n, n')$ with Boolean variables $d_1, \dots, d_n, d'_1, \dots, d'_{n'}$ such that $C_R(d_1, \dots, d_n, d'_1, \dots, d'_{n'}) = 1$ iff $((d_1, \dots, d_n), (d'_1, \dots, d'_{n'})) \in R$. For more about succinct circuit representation, see [17, 4, 36].

Problem 5.5. (TDP/PDP for symbolic relations) We are given a tuple $I = (n_{\mathcal{I}}, n_{\mathcal{O}}, n_{\mathcal{B}}, C_R)$ where $n_{\mathcal{I}}, n_{\mathcal{O}}, n_{\mathcal{B}}$ are logarithmic in the size of some domains $\mathcal{I}, \mathcal{O}, \mathcal{B}$ respectively and C_R is a circuit that describes a relation $R \subseteq \mathcal{I} \times \mathcal{O}$. The problem is whether there exist circuits C_{R_1}, C_{R_2} that describe relations $R_1 \subseteq \mathcal{I} \times \mathcal{B}$, and $R_2 \subseteq \mathcal{B} \times \mathcal{O}$ such that (R_1, R_2) meet the TDP/PDP conditions.

For TDP/PDP, as in Claim 5.2, the problem becomes trivial when $n_{\mathcal{B}} \geq \min\{n_{\mathcal{I}}, n_{\mathcal{O}}\}$. Note that a variant of PDP, in which the required relations in the solution are functions can be viewed as an instance of the problem of Boolean functional synthesis, e.g. [16, 9].

We next show that TDP/PDP are NEXPTIME-complete. We obtain this result by applying the computational-complexity theory of succinct-circuit representations for "logtime" reductions [4] to the NP-hardness reductions described in Section 5.1. We describe this in details.

A reduction from languages $A \subseteq \Sigma^*$ to $B \subseteq \Sigma^*$, for a finite alphabet Σ , is a function $f : A \rightarrow B$ such that $x \in A$ iff $f(x) \in B$. The function f is called a *logtime reduction* if the i -th symbol of $f(x)$ can be computed in a time logarithmic of the size of x . This is done by using a so called "direct-input-access" Turing machine, which has a specific "index" tape in which a binary index i is written and then the i -th symbol of the input string x is computed. See [4], which also states a generalization of the following:

Theorem 5.6. (Balcazar, Lozano, Toran [4]) *For every language $B \subseteq \Sigma^*$, if B is NP-hard under logtime reducibility, then the circuit representation of B is NEXPTIME-hard under polynomial-time reducibility.*

Since "standard" NP-hard problems (under polynomial-time reducibility) tend to be NP-hard also under logtime reducibility, the crux of Theorem 5.6 is that a succinct representation of a given problem in a form of a circuit does not necessarily ease the complexity of the solving the problem.

From Theorem 5.6 we get:

Theorem 5.7. *TDP/PDP for symbolic relations are NEXPTIME-complete.*

Proof. Membership in NEXPTIME is easily shown since the explicit variant can be obtained from the circuit representation in exponential time and both explicit TDP/PDP are in NP. For hardness We show that each of the reduction in a chain reductions from SAT to either GT18 or SP5 (see Section 5.1), and then to TDP or PDP, is a logtime reduction. For languages A and B , we use the notations $A \leq_m^P B$ to denote that A is polynomial-time reduced to B , and $A \leq^{lt} B$ to denote that A is logtime reduced to B (indices in brackets of the language indicate problem number in [18]). We have that $SAT(LO1) \leq_m^P 3SAT(LO2) \leq_m^P KCOLORABILITY(GT4) \leq_m^P PARTITIONCLIQUE(GT15) \leq_m^P CCBS(GT18) \leq_m^P TDP$, where all the reductions are described in [23] apart from the last reduction which is Theorem 5.4 from this paper, and the one before last which can be found in [21]. Similarly, we have that $SAT(LO1) \leq_m^P CLIQUE(GT19) \leq_m^P VC(GT1) \leq_m^P SC(SP5) \leq_m^P PDP$ where all the reductions are described in [23], the last reduction is again from Theorem 5.4.

As every problem A is a set of binary strings (words) x , and as every reduction from A to B is a function f such that $x \in A$ iff $f(x) \in B$, we have that the reductions described above are standard in the sense that only a few bits are required from x in order to determine the identity of the j 'th bit of $f(x)$. For example, in the reduction from $3SAT(LO1)$ to $CLIQUE(GT19)$, the output is a binary string that represents an adjacency matrix of a graph of size polynomial to the size of the input (and this size can be found in logarithmic time, see [4]). The content of every cell (i, j) of that matrix can be determined by whether two satisfying assignments to $3SAT$, one for a certain clause c and one for a certain clause c' in the input share the same variable (see [23] for more details). In this reduction, the output formula is standard, in the sense that the location of the clauses c, c' in the input string is easily found. As such, determining a certain output bit requires only a small (logarithmic) access to the input string, followed by a small (also logarithmic) comparison of the obtained input process. Similarly, the Cook-Levin reduction from any NP problem to SAT is standard in that sense as well, and therefore the output as the SAT formula can be obtained from logarithmic input and comparison. See [18] for more details.

To show that the reductions from $CCBS$ to TDP and from SC to PDP are logtime reduction, we need to be more accurate in the encoding of TDP/PDP. Indeed we can encode all problems as an adjacency matrix of the relation $R \subseteq \mathcal{I} \times \mathcal{O}$ in addition to $k \geq 0$ (in unary) that describes the size of the intermediate domain $\mathcal{B}$. The input for $CCBS$ is encoded as an adjacency matrix and some $k \geq 0$. A standard encoding for SC is also as an adjacency matrix (where the cell $(i, j) = 1$ iff the element a_i is a member of the set c_j) and some $k \geq 0$. As such, these reductions leave the input as is, since the difference between the instances of TDP/PDP, CCBS and SC is only in the semantics.

All in all we have that (i) SAT is NP-hard under logtime reductions, and (ii) every reduction described above is a logtime reduction. Logtime reductions are not transitive, however by using the Conversion Lemma from [4] we have that this chain of reduction for the succinct versions of the languages holds under polynomial-time reduction. Therefore by transitivity of polynomial-time reductions we get that TDP/PDP is NEXPTIME hard.

As a side note, see that the definitions of logtime reductions and the Conversion Lemma from [4] can be easily extended to poly-logtime reduction. Such reductions do preserve transitivity. $\square$

5.3. Automatic relations. In many applications we need to consider input and output as streams of symbols, with some desired relation between the input stream and the output stream. The most basic description for such systems, the one that we explore in this work, is when the domains are (possibly infinite) sets of finite words over finite alphabets, and the given relation is a regular language, given as a deterministic finite automaton (DFA), over the product alphabet of the input and output domains. The setting that we consider in this paper is of *automatic relations*. Automatic relations provide a context for a rich theory of automatic structures, cf. [24, 41], with a solvable decision for first-order logic. We follow here the convention in *regular model checking*, cf. [8], where *length-preserving* automatic relations, with input and output symbols interleaved, are used as input/output specifications for each step of reactive systems. (Thus, unlike the definitions in [24, 41], we do not allow padding.)

Given finite alphabets Σ, Σ' with domains $D \subseteq \Sigma^*, D' \subseteq \Sigma'^*$, and a relation $R \subseteq D \times D'$, we say that a DFA $\mathcal{A}_R$ over alphabet $(\Sigma \times \Sigma')$ describes R if: $(\vec{d}, \vec{d}') \in L(\mathcal{A}_R)$ if and only if $(\vec{d}, \vec{d}') \in R$. Note the slight abuse of notation as $L(\mathcal{A}_R)$ describes words in $(\Sigma \times \Sigma')^*$ while R describes words in $\Sigma^* \times \Sigma'^*$. Thus, we assume that input and output streams have the same lengths, that is, $(\vec{d}, \vec{d}') \in R$ implies that $|\vec{d}| = |\vec{d}'|$.

Problem 5.8. (TDP/PDP for automatic relations) We are given a tuple $I = (\Sigma_{\mathcal{I}}, \Sigma_{\mathcal{B}}, \Sigma_{\mathcal{O}}, \mathcal{A}_R)$, where $\Sigma_{\mathcal{I}}, \Sigma_{\mathcal{B}}$, and $\Sigma_{\mathcal{O}}$ are finite alphabets, and $\mathcal{A}_R$ is a DFA that describes a relation $R \subseteq \Sigma_{\mathcal{I}}^* \times \Sigma_{\mathcal{O}}^*$. The problem is whether there exist DFAs $\mathcal{A}_{R_1}, \mathcal{A}_{R_2}$ that describe relations $R_1 \subseteq \Sigma_{\mathcal{I}}^* \times \Sigma_{\mathcal{B}}^*$ and $R_2 \subseteq \Sigma_{\mathcal{B}}^* \times \Sigma_{\mathcal{O}}^*$ such that (R_1, R_2) meet the TDP/PDP conditions.

As in the explicit and symbolic cases, the problem becomes non trivial for TDP/PDP only when $|\Sigma_{\mathcal{B}}| < \min\{|\Sigma_{\mathcal{I}}|, |\Sigma_{\mathcal{O}}|\}$. While TDP/PDP in the symbolic setting can be solved by reduction to the explicit setting, this cannot be done here as the domains are possibly infinite. Indeed, automatic-relation TDP/PDP seems to be a challenging problem. We next conjecture that automatic-relation TDP/PDP are undecidable and explain the motivation for this conjecture and why an automatic-theoretic approach may not be helpful for TDP/PDP. We then show that even for the most basic case, in which the given relation is the equality relation, TDP/PDP can already be viewed as an algorithmic problem in automata that is equivalent to the Positivity Problem whose decidability is still open [42, 34]. We next show that for intermediate alphabet that is of exponential size, TDP/PDP on automatic relations can be reduced to TDP/PDP on binary intermediate alphabet. Finally we show by an automata-theoretic approach that a “strategic” variant of PDP, in which the required relations are in form of transducers, is decidable, and in fact is in EXPTIME.

5.3.1. An undecidability conjecture for TDP/PDP. A notable positive result about automatic relations is the decidability of their first-order theories [24, 41]. On the other hand, most second-order problems over automatic relations are undecidable; for example, checking the existence of an Eulerian cycle in an automatic graph is highly undecidable [29]. TDP/PDP are essentially second-order problems—we ask for the existence of R_1 and R_2 under the TDP/PDP conditions. Our conjecture is that this problem is undecidable. We provide here intuition to justify this conjecture.

We show in Section 6 below that TDP/PDP are decidable for automatic relations when a hint, in the form of a DFA for R_1 or R_2 is given. The idea is that given one component, say R_1 , there is a maximal second component R_2 that can be obtained from R_1 so that if R_1 is the first component of some proper decomposition, then R_2 is a proper second component.

When R_1 is given as a DFA, we can then construct a DFA for R_2 . Can we leverage this idea towards solving the problem in full? That is, can we search for, say, an automaton for R_1 that together with the maximal R_2 constructed below, forms a solution for TDP/PDP?

It is tempting to try to use an automata-theoretic approach similar to the strategic PDP (see Section 5.4 below) in which we consider representing $R_1 \subseteq (\Sigma_I \times \Sigma_B)^*$ as a labeled tree $\tau_1 : (\Sigma_I \times \Sigma_B)^* \mapsto \{0, 1\}$. Then we can try to define a tree automaton $\mathcal{A}_1$ that accepts a tree τ_1 iff it is a correct hint for the total or partial decomposition of an input/output relation $R \subseteq (\Sigma_I \times \Sigma_B)^*$. The difficulty is that such an automaton has to check two properties of τ_1 : (1) The domain of R_1 has to be equal to the domain of R . This is essentially a requirement on the projection of τ_1 on Σ_I . (2) The composition of R_1 with the maximal R_2 is contained in or equal to R . This is essentially a requirement on the projection of τ_1 on Σ_B . Known automata-theoretic techniques exist for dealing with projection of trees, see for example [26]. We currently, however, have no known technique to deal with two orthogonal projections, as we have here.

Of course, this argument only shows that a particular technique to attack the problem is unlikely to be successful. To justify the conjecture, we note that the dual-projection problem is reminiscent to the problem of distributed temporal synthesis, which was shown to be undecidable [39] (though there is no obvious formal connection between the decomposition problem and the distributed-synthesis problem). There we are representing an overall system strategy as a tree where tree labels correspond to actions of system components and tree edges correspond to environment actions. When different systems components are expected to act independently, without knowledge of actions by other components, overall systems strategies has to be decomposed into separate strategies as projections of the system strategy. The similarity between the two problems motivates us to formulate the following conjecture.

Conjecture 5.9. TDP/PDP for automatic relations is undecidable.

5.3.2. Decomposing equality. To support the claim of how non-trivial the decomposition problem is, we consider a more simple and abstract variant of automatic TDP. We consider a very simple case in which the given automatic relation is trivial, and the intermediate alphabet is binary. We do not even require the decomposed relation R_1, R_2 to be realized by automata, although we do require the length of every matching words in R_1 and R_2 to be the same. Specifically, given a regular language L over Σ_I^* , let $R_L^- \subseteq \Sigma_I^* \times \Sigma_I^*$ be the *equality* relation over L . i.e. $R = \{(w, w) | w \in L\}$. Given L over Σ_I^* and $\Sigma_B = \{0, 1\}$, we ask whether there are relations (R_1, R_2) that meet the TDP properties with respect to R_L^- .

First see that we can assume wlog that (R_1, R_2) are functions since existence of such relations leads to existence of such functions. Next, note that R_1 cannot relate two distinct words in L to the same word in Σ_B^* . This is because otherwise a word from Σ_B^* has to meet two distinct Σ_I^* elements, thus either break the equality property or break the TDP property. Therefore we have that R_1 is an injection from Σ_I^* to Σ_B^* . As such, finding an R_1 that is such an injection also gives us $R_2 = R_1^{-1}$.

Let L_n be the words in L of size n . Note that if there is n for which $|L_n| > 2^n$ then no such R_1 can be found. If, however, for every n we have $|L_n| \leq 2^n$ then a function R_1 can be simply realized by ordering the words in every L_n in lexicographic order and relating each one to the Σ_B^n word that encodes the index in binary.

Therefore the problem of TDP in this setting is reduced to the following problem: given a regular language L over a finite alphabet Σ , where for every n , L_n is the set of words in L of size n , does $|L_n| \leq 2^n$ for every n ? We call this problem the *The Exponential-Bound Problem* (EBP). The following theorem, however, shows that EBP is equivalent to the problem of *Positivity*, described below, whose decidability has been famously open for decades [42, 34]. Although this is not a direct reduction to automatic TDP, this relation indicates to the hardness of solving automatic TDP on even a simple relation.

A *linear recurrence sequence* (LRS) is a sequence of integers $\langle u_n \rangle_{n=0}^\infty$ satisfying a recurrence relation: there exist constants integers $a_1, a_2, \dots, a_k$ such that, for all $n \geq 0$, $u_{n+k} = a_1 u_{n+k-1} + a_2 u_{n+k-2} + \dots + a_k u_n$. If the initial values $u_0, \dots, u_{k-1}$ of the sequence are provided, the recurrence relation defines the rest of the sequence uniquely. Given a linear recurrence sequence (LRS) $\langle u_n \rangle_{n=0}^\infty$, the *Positivity Problem* asks whether all terms of the sequence are non-negative.

Theorem 5.10. *EBP is Equivalent to Positivity.*

Proof. We first show that Positivity reduces to EBP. Let $\langle u_n \rangle_{n=0}^\infty$ be an LRS; we show how positivity for the sequence $\langle -u_n \rangle_{n=0}^\infty$ can be formulated as an EBP problem. To this end, we invoke Corollary 4 from [2] to obtain a rational stochastic matrix M of size $c \times c$ (denoted $\tilde{Q}$ in the proof of Proposition 2 from *op. cit.*, where the value of c is also stated) such that, for all $n \geq 0$,

$$(M^{2n+1})_{1,2} \leq \frac{1}{4} \quad \text{iff} \quad u_n \leq 0 \quad \text{iff} \quad -u_n \geq 0, \quad (5.1)$$

in other words the positivity of $\langle -u_n \rangle_{n=0}^\infty$ is violated iff there is some n such that the $(1, 2)$ -th entry of M^{2n+1} is strictly larger than $1/4$.

In fact, as noted in the comments following Corollary 4 in [2], M can be chosen so that its entries are dyadic rationals, i.e., having denominator some power of 2. Let us therefore assume this to be the case, and let 2^p be the largest such power. Write $J = 2^p M$ and $N = (2^p M)^2$. Then J and N are square matrices with non-negative integer coefficients, and hence there is some DFA $\mathcal{A}$ such that $(J \cdot N^n)_{1,2}$ is the number of words of length $n+1$ accepted by $\mathcal{A}$. More precisely, $\mathcal{A}$ has initial state s , and c further states $q_1, \dots, q_c$. The single accepting state is q_2 . To define the transition function, if the $(1, j)$ -th entry of J is ℓ , then we postulate ℓ transitions going from state s to state q_j , each labelled with a new (fresh) letter. Likewise, if the (i, j) -th entry of N is ℓ , then we include ℓ transitions going from q_i to q_j , again for each one using a new letter as label. In this way, J and N can be viewed as the adjacency matrices of the underlying directed multigraph of $\mathcal{A}$, and $(J \cdot N^n)_{1,2}$ counts the number of paths in $\mathcal{A}$ going from s to q_2 in $n+1$ steps. Since by construction, different paths give rise to different words, $(J \cdot N^n)_{1,2}$ does indeed correspond to the number of words of length $n+1$ accepted by $\mathcal{A}$.

Writing $L(\mathcal{A}) = L$, Eq. (5.1) becomes, for all $n \geq 0$,

$$L_{n+1} \leq \frac{2^p 2^{2pn}}{4} = \frac{2^{2p(n+1)}}{2^{p+2}} \quad \text{iff} \quad -u_n \geq 0. \quad (5.2)$$

We now modify automaton $\mathcal{A}$ by lengthening every transition in $\mathcal{A}$ by a factor of $2p$; more precisely, for every transition $q \rightarrow q'$, create $2p-1$ fresh non-accepting states $r_1, \dots, r_{2p-1}$ and replace $q \rightarrow q'$ by the sequence $q \rightarrow r_1 \rightarrow r_2 \rightarrow \dots \rightarrow r_{2p-1} \rightarrow q'$, all labelled with the same letter as the original transition. The initial and accepting states otherwise remain unchanged. Let us call the resulting DFA $\mathcal{A}'$, with accepted language

$L(\mathcal{A}') = L'$. In moving from $\mathcal{A}$ to $\mathcal{A}'$, the net effect has been to increase the length of every accepted word by a factor of $2p$; note also that if m is not a multiple of $2p$, then $L'_m = 0$.

Combining the above with Eq. (5.2), we conclude that the LRS $\langle -u_n \rangle_{n=0}^\infty$ is positive iff for all $m \geq 0$, $L'_m \leq \frac{2^m}{2^{p+2}}$, i.e., $2^{p+2}L'_m \leq 2^m$.

Inflating the alphabet size of $\mathcal{A}'$ by a factor of 2^{p+2} , we can easily manufacture a DFA $\mathcal{A}''$ with accepted language $L(\mathcal{A}'') = L''$ having the property that, for all $m \geq 0$, $L''_m = 2^{p+2}L'_m$. It therefore follows that the LRS $\langle -u_n \rangle_{n=0}^\infty$ is positive iff for all $m \geq 0$, $L''_m \leq 2^m$, which completes the reduction of Positivity to EBP.

Finally, since the series obtained from the number of distinct words of length n in an automaton makes an LRS [35], we have that EBP also reduces to Positivity, so that the two problems are in fact equivalent. $\square$

5.3.3. Reduction to binary alphabet. Since the size of the intermediate alphabet plays a crucial role in the solution of TDP/PDP, one may ask whether it suffices to search for solutions for only the binary case. We show that the answer is positive for intermediate alphabet that is of size of exponent. Specifically we show by reduction that for every automatic TDP/PDP instance $I = (\Sigma_{\mathcal{I}}, \Sigma_{\mathcal{B}}, \Sigma_{\mathcal{O}}, \mathcal{A}_R)$, where $|\Sigma_{\mathcal{B}}| = 2^m$ for some $m > 0$, there is a TDP/PDP instance $I' = (\Sigma_{\mathcal{I}}, \{0, 1\}, \Sigma_{\mathcal{O}}, \mathcal{A}_{R'})$ for some relation R' such that I has a solution if and only if I' has a solution.

To show our reduction, we first need to reason over automatic decomposition in the regular expressions representation level. Therefore we give below the following lemma that characterizes a possible TDP/PDP solution as regular expressions.

Lemma 5.11. *A TDP/PDP for automatic relations instance $I = (\Sigma_{\mathcal{I}}, \Sigma_{\mathcal{B}}, \Sigma_{\mathcal{O}}, \mathcal{A}_R)$ has a solution if and only if there is a regular expression $S \subseteq (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})^*$ such that the following happens:*

- (1) *(Only for the TDP variant) For every $(\vec{i}, \vec{o}) \in R$ there is $\vec{b} \in \Sigma_{\mathcal{B}}^*$ such that $(\vec{i}, \vec{b}, \vec{o}) \in L(S)$.*
- (2) *(Only for the PDP variant) For every $\vec{i} \in \text{Dom}(R)$ there are $\vec{b} \in \Sigma_{\mathcal{B}}^*$, and $\vec{o} \in \Sigma_{\mathcal{O}}^*$ such that $(\vec{i}, \vec{b}, \vec{o}) \in L(S)$.*
- (3) *(For both variants) For every $(\vec{i}, \vec{b}, \vec{o}), (\vec{i}', \vec{b}', \vec{o}') \in L(S)$, if $\vec{b} = \vec{b}'$ then $(\vec{i}, \vec{o}'), (\vec{i}', \vec{o}) \in R$.*

Proof. Assume the instance I has a TDP/PDP solution $(\mathcal{A}_{R_1}, \mathcal{A}_{R_2})$. Let $\mathcal{A}'$ be the product automata of $\mathcal{A}_1$ and $\mathcal{A}_2$ followed by eliminating the edges (i, b, b', o) where $b \neq b'$, and then eliminating the first alphabet of $\Sigma_{\mathcal{B}}$. Therefore $\mathcal{A}'$ is over the alphabet $(\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})$. Let S be the regular expression that describes $\mathcal{A}'$. First verify (3): let $(\vec{i}, \vec{b}, \vec{o}), (\vec{i}', \vec{b}, \vec{o}') \in L(S)$. Then $(\vec{i}, \vec{b}), (\vec{i}', \vec{b}) \in R_1$ and $(\vec{b}, \vec{o}), (\vec{b}, \vec{o}') \in R_2$ which makes $(\vec{i}, \vec{o}'), (\vec{i}', \vec{o}) \in R_1 \circ R_2$, hence both for TDP/PDP we have $(\vec{i}, \vec{o}'), (\vec{i}', \vec{o}) \in R$. For TDP we verify (1): Let $(\vec{i}, \vec{o}) \in R$, then there is $\vec{b}$ such that $(\vec{i}, \vec{b}) \in R_1$ and $(\vec{b}, \vec{o}) \in R_2$ therefore $(\vec{i}, \vec{b}, \vec{o}) \in L(S)$. For PDP we verify (2): let $\vec{i} \in \text{Dom}(R)$, then there is $\vec{o}$ such that $(\vec{i}, \vec{o}) \in R$. Then as before there is $\vec{b}$ such that $(\vec{i}, \vec{b}) \in R_1$ and $(\vec{b}, \vec{o}) \in R_2$ therefore $(\vec{i}, \vec{b}, \vec{o}) \in L(S)$.

Next, assume S is a regular expression over $(\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})^*$ that meets conditions (1)-(3) for the TDP/PDP instance I . Let $\mathcal{A}'$ be the automaton that describes S . Let $\mathcal{A}_1$ be the automaton obtained from $\mathcal{A}'$ by eliminating the $\Sigma_{\mathcal{O}}$ letters and determinizing. Same let $\mathcal{A}_2$ be the automaton obtained from $\mathcal{A}'$ by eliminating the $\Sigma_{\mathcal{I}}$ letters and determinizing. Let R_1, R_2 be the regular relations that $\mathcal{A}_1, \mathcal{A}_2$ respectively describe. We see that (R_1, R_2) meet

the TDP/PDP conditions. Let $\vec{b} \in \text{Img}(R_1)$. Then there is $\vec{i}$ such that $(\vec{i}, \vec{b})$ is an accepting word for $\mathcal{A}_1$, therefore there exists $\vec{o}$ such that $(\vec{i}, \vec{b}, \vec{o})$ is accepting for $\mathcal{A}'$. Therefore $(\vec{b}, \vec{o})$ is accepting for $\mathcal{A}_2$ so $\vec{b} \in \text{Dom}(R_2)$. For TDP let $(\vec{i}, \vec{o}) \in R$. Then from condition (1) there is $\vec{b} \in \Sigma_{\mathcal{B}}^*$ such that $(\vec{i}, \vec{b}, \vec{o}) \in L(S)$, so $(\vec{i}, \vec{b}, \vec{o})$ is accepting for $\mathcal{A}'$, therefore $(\vec{i}, \vec{b}) \in R_1$, and $(\vec{b}, \vec{o}) \in R_2$ so $(\vec{i}, \vec{o}) \in R_1 \circ R_2$. For PDP, let $\vec{i} \in \text{Dom}(R)$. Then from condition (2) there are $\vec{b}, \vec{o}$ such that $(\vec{i}, \vec{b}, \vec{o}) \in L(S)$, so $(\vec{i}, \vec{b}, \vec{o})$ is accepting for $\mathcal{A}'$, therefore $(\vec{i}, \vec{b}) \in R_1$, and $(\vec{b}, \vec{o}) \in R_2$ so $\vec{i} \in \text{Dom}(R_1 \circ R_2)$. Finally for both TDP/PDP assume $(\vec{i}, \vec{o}) \in R_1 \circ R_2$. Then there is $\vec{b}$ such that $(\vec{i}, \vec{b}) \in R_1$, and $(\vec{b}, \vec{o}) \in R_2$. Then there are $\vec{i}', \vec{o}'$ such that $(\vec{i}, \vec{b}, \vec{o}')$ and $(\vec{i}', \vec{b}, \vec{o})$ are accepting for $\mathcal{A}'$ and therefore in $L(S)$. Then from condition (3), $(\vec{i}, \vec{o}) \in R$. $\square$

Using Lemma 5.11, we prove the following theorem.

Theorem 5.12. *There is a TDP/PDP solution to $I = (\Sigma_{\mathcal{I}}, \Sigma_{\mathcal{B}}, \Sigma_{\mathcal{O}}, \mathcal{A}_R)$ where $|\Sigma_{\mathcal{B}}| = 2^m$ for some $m > 0$, if and only if there is a TDP/PDP solution to $I' = (\Sigma_{\mathcal{I}}, \{0, 1\}, \Sigma_{\mathcal{O}}, \mathcal{A}_{R'})$.*

The idea behind the reduction is to encode every letter in $\Sigma_{\mathcal{B}}$ in binary, and use this encoding for the construction of I' . $\mathcal{A}_{R'}$ is therefore constructed from $\mathcal{A}_R$ by replacing every edge labeled (i, o) with a path of m edges, each with the same label (i, o) .

Proof. Let bin be a function in which $\text{bin}(b)$ is a binary encoding of the letter $b \in \Sigma_{\mathcal{B}}$ in m bits, and let $\text{bin}_j(b)$ be the j 'th bit in $\text{bin}(b)$. Since $\Sigma_{\mathcal{B}}$ is of size 2^m , the function bin is a bijection. Assume that I has a solution $(\mathcal{A}_{R_1}, \mathcal{A}_{R_2})$. We first construct an automaton $\mathcal{A}_{R'_1}$ over alphabet $(\Sigma_{\mathcal{I}} \times \{0, 1\})$ from $\mathcal{A}_{R_1}$ by replacing every edge labeled (i, b) with an m -edges path with edge labels $(i, \text{bin}_0(b)), \dots, (i, \text{bin}_{m-1}(b))$, thus such a path describes the word $(i^m, \text{bin}(b))$ (where i^m is the letter i concatenated m times). Same, we construct an automaton $\mathcal{A}_{R'_2}$ over alphabet $(\{0, 1\} \times \Sigma_{\mathcal{O}})$ from $\mathcal{A}_{R_2}$ by replacing every edge labeled (b, o) with an m -edges path with edge labels that describes the word $(\text{bin}(b), o^m)$. Then since (R_1, R_2) solve TDP/PDP for the instance I , we have that (R'_1, R'_2) solve TDP/PDP for the instance I' .

The other side of the proof is more involved since suppose we assume $(\mathcal{A}_{R'_1}, \mathcal{A}_{R'_2})$ solve I' . Then we cannot guarantee that these automata have an "m-path" structure that can reverse the construction we have just described. For that, we define by induction an homomorphism $h : (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})^* \rightarrow (\Sigma_{\mathcal{I}}^m \times \{0, 1\}^m \times \Sigma_{\mathcal{O}}^m)^*$ as follows. for every letter (i, b, o) , we define $h((i, b, o)) = (i^m, \text{bin}(b), o^m)$. For a word $\sigma \in (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})^*$ and a letter $x \in (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})$ we define $h(\sigma x) = h(\sigma)h(x)$. By definition, h is an homomorphism, note that h is also an injection.

Now assume that there is a TDP/PDP solution to $I' = (\Sigma_{\mathcal{I}}, \{0, 1\}, \Sigma_{\mathcal{O}}, \mathcal{A}_{R'})$, and recall that $\mathcal{A}_{R'}$ is obtained from $\mathcal{A}_R$ by replacing every letter (i, o) with a word (i^m, o^m) . Let S' be the regular expression obtained from Lemma 5.11 w.r.t. I' , and let $L(S')$ be the regular language of S' . Since every word in $L(S')$ is a concatenation of strings of the form (i^m, t, o^m) for some $i \in \Sigma_{\mathcal{I}}$, $t \in \{0, 1\}^m$ and $o \in \Sigma_{\mathcal{O}}$, we have that $L(S') \subseteq (\Sigma_{\mathcal{I}}^m \times \{0, 1\}^m \times \Sigma_{\mathcal{O}}^m)^*$. Moreover, since bin is a bijection we have that for every such (i^m, t, o^m) , $(i, \text{bin}^{-1}(t), o)$ is well defined and $h(i, \text{bin}^{-1}(t), o) = (i^m, t, o^m)$, therefore the homomorphism h is onto $L(S')$. Since the pre-image under homomorphisms of a regular language is a regular language as well (see [25]), we have that there is a regular language $L(S)$ of a regular expression S over $(\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})^*$ such that $L(S) = h^{-1}(L(S'))$.

Since S' satisfies conditions (1)-(3) of Lemma 5.11 w.r.t the instance I' and since h is a bijection from $L(S)$ to $L(S')$, it is easy to check that S satisfies the conditions (1)-(3) of Lemma 5.11 w.r.t the instance I , which means that I has a TDP/PDP solution. $\square$

Theorem 5.12 is stated only for alphabet that are of exponent size. We do not yet know the solution for the general case. We thus have the following conjecture.

Conjecture 5.13. There is a reduction for TDP/PDP from an intermediate alphabet of arbitrary size to the binary alphabet.

Corollary 6.7 shows that TDP/PDP on automatic relations with unary intermediate alphabet is solvable.

5.4. Strategic automatic PDP. A specific version of PDP on automatic relation, that captures essential concepts in synthesis [27, 16, 39], is when we require the solution to be strategies. For that, we define *Strategic automatic PDP*, or Strategic PDP in short, as PDP on automatic relations in which the required relations R_1 and R_2 are functions (defined in Section 3) that are represented by transducers T_1 and T_2 , respectively.

Strategic PDP can be viewed as a game of incomplete information. Since the information “flows” in one direction, however, the key to proving decidability for this problem is to view the problem as a one-way chain communication of distributed synthesis from [27] to synthesize the required transducers.

Theorem 5.14. *Strategic PDP is in EXPTIME.*

Proof. The proof for Theorem 5.14 lies heavily on definitions and terminology from [27]. We first describe the intuition of the proof by providing a high-level outline. Note that T_1 is a finite-state realization of a function $f_1 : \Sigma_I^* \rightarrow \Sigma_B$, and T_2 is a finite-state realization of a function $f_2 : \Sigma_B^* \rightarrow \Sigma_O$. We first consider trees of the form $\tau_1 : \Sigma_I^* \rightarrow \Sigma_B \times \Sigma_O$. Such a tree represents simultaneously both f_1 and f_2 . A branch of this tree can be viewed as a word $w \in (\Sigma_I \times \Sigma_B \times \Sigma_O)^\omega$. By running the DFA $\mathcal{A}_R$ on w we can check that every prefix of w is consistent with R . Thus, we can construct a deterministic automaton $\mathcal{A}_R^t$ on infinite trees that checks that all prefixes of all branches in τ_1 are consistent with R .

Note, however, that in τ_1 the Σ_O values depend not only on the Σ_B values but also on the Σ_I values, while in f_2 the domain is Σ_B^* . So now we consider a tree $\tau_2 : \Sigma_B^* \rightarrow \Sigma_O$. We now simulate $\mathcal{A}_R^t$ on τ_2 (cf., [27]). The idea is to match each branch of τ_1 with a branch of τ_2 according to the $\mathcal{B}$ -values. This means that we have several branches of the run of $\mathcal{A}_R^t$ running on one branch of τ_2 . We thus obtain an alternating tree automaton $\mathcal{A}_2$, whose size is linear in the size of $\mathcal{A}_R$. We can now check in exponential time non-emptiness of $\mathcal{A}_2$, to see if a function f_2 exists. If $\mathcal{A}_2$ is non-empty, then we obtain a witness transducer T_2 whose size is exponential in $\mathcal{A}_R$.

Finally, we consider a tree $\tau'_1 : \Sigma_I^* \rightarrow \Sigma_B$ that represents f_1 . We run both $\mathcal{A}_R$ and T_2 on each branch of τ'_1 , where T_2 generates the $\mathcal{O}$ -values that were present in τ_1 but not in τ'_1 . Using these values $\mathcal{A}_R$ checks that each prefix is consistent with R . We have obtained a deterministic tree automaton $\mathcal{A}_1$, whose size is exponential in $\mathcal{A}_R$. We can now check non-emptiness of $\mathcal{A}_1$, and obtain a witness transducer T_1 . Thus, we can solve Strategic PDP for automatic relations in exponential time.

We now describe the proof of Theorem 5.14 in detail. An (alternating) tree-automaton that runs over Σ -labeled Γ -tree is a tuple $\mathcal{A} = (\Gamma, \Sigma, Q, q_0, \delta, \alpha)$ where Σ, Γ are finite alphabets, Q is a finite set of states, q_0 is an initial state, $\alpha \subseteq Q^\omega$ is an acceptance condition, and $\delta : Q \times \Sigma \rightarrow B^+(\Gamma \times Q)$ is the transition function, where $B^+(\Gamma \times Q)$ is the set of positive Boolean formulas over $\Gamma \times Q$. See more details in [27]. To simplify the notation we denote the elements of an automaton by the $\cdot$ sign. For example $\mathcal{A}.Q$ denotes the set of states Q of $\mathcal{A}$.

By using the *xray* definition together with Theorem 4.2 from [27], we can assume without loss of generality that $\mathcal{A}$ runs over $(\Gamma \times \Sigma)$ -labeled Γ -trees. That is, the label contains information about the direction of the tree as well. Therefore when we consider trees of the form $\Sigma_{\mathcal{I}}^* \rightarrow \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}}$, we can assume a tree $\Sigma_{\mathcal{I}}^* \rightarrow \Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}}$, thus a branch of such tree can be viewed as a word $w \in (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})^\omega$.

Recall that R is the given relation for the Strategic PDP problem. We now use the DFA of R , $\mathcal{A}_R = (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{O}}, Q, q_0, \delta, F)$, to construct a deterministic tree automaton $\mathcal{A}_R^t$ over trees $\tau : \Sigma_{\mathcal{I}}^* \rightarrow \Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}}$ that checks that all prefixes of all branches in any accepting tree τ are consistent with R .

We now define $\mathcal{A}_R^t$ as follows. $\mathcal{A}_R^t = (\Sigma_{\mathcal{I}}, \Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}}, Q, q_0, \delta, \alpha)$ as a $(\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})$ -labeled $\Sigma_{\mathcal{I}}$ -tree automata. in which $\mathcal{A}_R^t.Q = \mathcal{A}_R.Q$, $\mathcal{A}_R^t.q_0 = \mathcal{A}_R.q_0$, $\mathcal{A}_R^t.\alpha = \mathcal{A}_R.Q$. Finally we define

$$\mathcal{A}_R^t.\delta(q, (i, b, o)) = \bigwedge_{i' \in \Sigma_{\mathcal{I}}} (i', \mathcal{A}_R.\delta(q, (i, o))) \quad (5.3)$$

See that $\mathcal{A}_R^t$ is a deterministic (note the slight unusual notation) tree automaton that checks that all prefixes of all branches are consistent with R .

From $\mathcal{A}_R^t$ we obtain an automaton $\mathcal{A}_R^t = \text{cover}(\mathcal{A}_R^t)$ that accepts an $(\Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})$ -label $\Sigma_{\mathcal{I}}$ -tree τ iff $\mathcal{A}_R^t$ accepts *xray*(τ), see [27] for more details.

We next use $\mathcal{A}_R^t$ to construct the Strategic PDP solution transducers T_1, T_2 , starting with T_2 , that describe trees τ_1, τ_2 respectively. To construct τ_1 we note that in τ_1 the $\Sigma_{\mathcal{O}}$ values depend not only on the $\Sigma_{\mathcal{B}}$ values but also on the $\Sigma_{\mathcal{I}}$ values, while in τ_2 the domain is $\Sigma_{\mathcal{B}}^*$. We construct τ_2 by obtaining an alternating tree automaton $\mathcal{A}_2$, which size is linear in the size of $\mathcal{A}_R$. Every tree in $L(\mathcal{A}_2)$ can be a partial solution for Strategic PDP.

To obtain $\mathcal{A}_2$ we use definitions from [27] with a small modification of the construction: we do not need the *delay* property used in [27]. We first look at $(\Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})$ -label $\Sigma_{\mathcal{I}}$ -trees $(\Sigma_{\mathcal{I}}^*, f)$, where $f : \Sigma_{\mathcal{I}}^* \rightarrow (\Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}})$ is the tree labeling function, appropriately defined in [27]. We say that f is a *tree composition* of a $\Sigma_{\mathcal{B}}$ -label $\Sigma_{\mathcal{I}}$ -trees $(\Sigma_{\mathcal{I}}^*, f_1)$ and a $\Sigma_{\mathcal{O}}$ -label $\Sigma_{\mathcal{B}}$ -trees $(\Sigma_{\mathcal{B}}^*, f_2)$ (denoted $f = f_1 + f_2$) if for every $i_1, i_2 \in \Sigma_{\mathcal{I}}$ and every word $\sigma \in \Sigma_{\mathcal{I}}^*$ we have:

- $f(\epsilon) = f_1(\epsilon) \times f_2(f_1(\epsilon))$.
- $f(\sigma) = f_1(\sigma) \times f_2(f_1'(\sigma))$.

Where $f_1' : \Sigma_{\mathcal{I}}^* \rightarrow \Sigma_{\mathcal{B}}^*$ is a function that preserves the *memory* of f_1 . i.e. $f_1'(\epsilon) = f_1(\epsilon)$ and for every $i \in \Sigma_{\mathcal{I}}^*, \sigma \in \Sigma_{\mathcal{I}}^*$ we have $f_1'(\sigma \cdot i) = f_1'(\sigma) \cdot f_1(i)$.

Next, for a set T of $(\Sigma_{\mathcal{O}} \times \Sigma_{\mathcal{B}})$ -labeled $\Sigma_{\mathcal{I}}$ trees, we say *shape _{$\mathcal{I}$}* (T) consists of all the $\Sigma_{\mathcal{O}}$ -labeled $\Sigma_{\mathcal{B}}$ trees $(\Sigma_{\mathcal{B}}^*, f_2)$ for which there exists a $\Sigma_{\mathcal{B}}$ -labeled $\Sigma_{\mathcal{I}}$ trees $(\Sigma_{\mathcal{I}}^*, f_1)$ such that the $(\Sigma_{\mathcal{O}} \times \Sigma_{\mathcal{B}})$ -labeled $\Sigma_{\mathcal{I}}$ tree $(\Sigma_{\mathcal{I}}^*, f_1 + f_2)$ is in T .

The following claim is analogical to Theorem 4.1 from [27].

Claim 5.15. Let $\mathcal{A}$ be a non-deterministic tree automaton over $(\Sigma_{\mathcal{O}} \times \Sigma_{\mathcal{B}})$ -labeled $\Sigma_{\mathcal{I}}$ trees. Then we can construct an alternating tree automaton $\mathcal{A}'$, denoted by $\text{shape}_{\mathcal{I}}(\mathcal{A})$ over $\Sigma_{\mathcal{O}}$ -labeled $\Sigma_{\mathcal{B}}$ trees such that $L(\mathcal{A}') = \text{shape}_{\mathcal{I}}(L(\mathcal{A}))$ and $\mathcal{A}, \mathcal{A}'$ have the same size and index.

Proof. We describe $\mathcal{A}'$. Although the construction is the same, this proof is different from [27] since we do not need *delay*.

Let $\mathcal{A} = (\Sigma_{\mathcal{I}}, \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}}, Q, q_0, \delta, \alpha)$. Then we have $\mathcal{A}' = (\Sigma_{\mathcal{B}}, \Sigma_{\mathcal{O}}, Q, q_0, \delta', \alpha)$, where for every $q \in Q, o \in \Sigma_{\mathcal{O}}$ we have

$$\delta'(q, o) = \bigvee_{b \in \Sigma_{\mathcal{B}}, \vec{s} \in \delta(q, (b, o))} (b, s_1) \wedge (b, s_2) \wedge \cdots \wedge (b, s_{|\Sigma_{\mathcal{I}}|}) \quad (5.4)$$

where $\vec{s} = (s_1, \dots, s_{|\Sigma_{\mathcal{I}}|})$.

That is, at every state q when reading a label o , $\mathcal{A}'$ guesses an element b from $\Sigma_{\mathcal{B}}$, and then "throws" the automaton in the directions that $\mathcal{A}$ would have taken with label (b, o) . Note that $\mathcal{A}'$ is an alternating automaton. Then for every $\Sigma_{\mathcal{O}}$ -labeled $\Sigma_{\mathcal{B}}$ tree τ we have that $\tau \in L(\mathcal{A}')$ iff $\tau \in \text{shape}_x(L(\mathcal{A}))$.

That concludes the proof for Claim 5.15.

Back to the construction, we define $\mathcal{A}_2 = \text{shape}_{\Sigma_{\mathcal{I}}}(\mathcal{A}'^t_R)$. We then construct $\mathcal{A}_2$ and check emptiness (in exponential time since $\mathcal{A}_2$ is an alternating automaton). If $\mathcal{A}_2$ is empty then we conclude that there is no Strategic PDP solution. Otherwise we get a tree $(\Sigma_{\mathcal{B}}^*, f_2)$ in a form of a transducer T_2 that is our partial Strategic PDP solution. We then use T_2 and $\mathcal{A}'^t_R$ to construct a tree automaton $\mathcal{A}_1$ over $\Sigma_{\mathcal{B}}$ -label $\Sigma_{\mathcal{I}}$ trees such that every solution of f_1 to $\mathcal{A}_1$ has the property that $f_1 + f_2 = f$. This is done as follows. Let $T_2 = (\Sigma_{\mathcal{B}}, \Sigma_{\mathcal{O}}, S, \nu, s_0, \text{Out})$ where $\nu : S \times \Sigma_{\mathcal{B}} \rightarrow S$ is the transition function and $\text{Out} : S \rightarrow \Sigma_{\mathcal{O}}$ is the output function of T_2 . Recall that $\mathcal{A}'^t_R = (\Sigma_{\mathcal{I}}, \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}}, Q, q'_0, \delta', \alpha')$. Then we construct an automaton $\mathcal{A}_1 = (\Sigma_{\mathcal{I}}, \Sigma_{\mathcal{B}}, Q' \times S, (q'_0, s_0), \delta_1, \alpha_1)$ where $\alpha_1 = \alpha' \times S$ (i.e. replace every element set $Q'' \subseteq Q'$ from α' with $Q'' \times S$), and δ_1 is defined as follows. For $q \in Q', s \in S$ and $b \in \Sigma_{\mathcal{B}}$, let $s' = \nu(s, b)$. Then we have:

$$\delta_1((q, s), b) = \bigvee_{\vec{z} \in \delta(q, (b, \text{Out}(s')))} ((z_1, s'), \dots, (z_{|\Sigma_{\mathcal{I}}|}, s')) \quad (5.5)$$

where $\vec{z} = (z_1, \dots, z_{|\Sigma_{\mathcal{I}}|})$.

That is, $\mathcal{A}_1$ simulates $\mathcal{A}'^t_R$, and in every state (q, s) when reads a label b , simulates T_2 appropriately to get the output label o obtained from the output of T_2 from state s and label b .

Note that $\mathcal{A}_1$ accepts the trees that match Claim 5.15 for $(\Sigma_{\mathcal{B}}^*, f_2)$ in the sense that $\mathcal{A}_1$ accepts exactly the trees $(\Sigma_{\mathcal{I}}^*, f_1)$ for which $f_1 + f_2 = f$. Finally we check emptiness for $\mathcal{A}_1$. Since $\mathcal{A}_2$ is non-empty, then from Claim 5.15 we have that $\mathcal{A}_1$ cannot be empty therefore we get a tree $(\Sigma_{\mathcal{I}}^*, f_1)$ in form of a transducer T_1 . This completes the construction. Since checking emptiness for alternating automata is in EXPTIME, the entire construction is in EXPTIME as well. $\square$

6. DECOMPOSITION WITH A HINT

Our results above indicate that fully automated decomposition is a hard problem. Can we ameliorate this difficulty by including a “human in the loop”? Indeed, in some decomposition scenarios, a part of a decomposition is already given, and the challenge is to find the complementary component. This can be thought of as a partial solution that is offered by a human intuition, e.g. a domain expert. In the context of our framework we have that a candidate to either R_1 or R_2 is given (in the relevant description formalism) as a “hint”. The question then is whether, given one possible component, a complementary component indeed exists, and can be constructed, such that both relations together meet the TDP/PDP conditions.

We discuss first TDP/PDP problems with a hint R_1 followed by a discussion with similar arguments for TDP/PDP problems with hint R_2 .

6.1. TDP/PDP with hint R_1 . TDP/PDP problem with a hint R_1 is formally defined as follows.

Definition 6.1. In the TDP/PDP problem with a hint R_1 we are given input, output and intermediate domain $\mathcal{I}, \mathcal{O}, \mathcal{B}$, relations $R \subseteq \mathcal{I} \times \mathcal{O}$ and $R_1 \subseteq \mathcal{I} \times \mathcal{B}$. The goal is to find whether there is a relation $R_2 \subseteq \mathcal{B} \times \mathcal{O}$ such that (R_1, R_2) meet the TDP/PDP conditions.

The exact nature of the domains and relations varies as before according to the problem setting (explicit, etc.). As we see, such a hint as a partial solution relaxes the computational difficulty of TDP/PDP, shown in previous sections, considerably. To that end, we show the following *maximum property* that is relevant for all TDP/PDP settings. Given a TDP/PDP instance with a hint R_1 , define a relation $R'_2 \subseteq (\mathcal{B} \times \mathcal{O})$ to be $R'_2 = \{(b, o) \mid \forall i \in \mathcal{I} ((i, b) \in R_1 \rightarrow (i, o) \in R)\}$. Note that $\text{Dom}(R'_2)$ can strictly contain $\text{Img}(R_1)$.

Lemma 6.2. *Every solution for TDP/PDP with a hint R_1 is contained in R'_2 and if there exists such a solution then R'_2 is a solution as well.*

Proof. We prove for TDP, we skip the proof for PDP that is almost identical. Let R_2 be a solution to TDP with a hint R_1 . We first see that $R_2 \subseteq R'_2$. Let $(b, o) \in R_2$. Suppose there exists i such that $(i, b) \in R_1$ and $(i, o) \notin R$. Then we have that $R_1 \circ R_2 \not\subseteq R$ which means that R_2 is not a solution, a contradiction. Therefore we have that for all i , if $(i, b) \in R_1$ then $(i, o) \in R$, hence $(b, o) \in R'_2$. To see that R'_2 is a solution, first see that since R_2 is a solution contained in R'_2 then $\text{Dom}(R_1 \circ R'_2) = \text{Dom}(R)$ and $\text{Img}(R_1) \subseteq \text{Dom}(R'_2)$. Let $(i, o) \in R_1 \circ R'_2$. Then there is a b such that $(i, b) \in R_1$ and $(b, o) \in R'_2$, which means that by the definition of R'_2 and since $(i, b) \in R_1$, it must be that $(i, o) \in R$. Finally, assume that $(i, o) \in R$. Then there is b such that $(i, b) \in R_1$, and $(b, o) \in R'_2$ for some solution R'_2 contained in R'_2 . Therefore $(b, o) \in R'_2$ so $(i, o) \in R_1 \circ R'_2$. $\square$

From Lemma 6.2 we get a simple method to solve TDP/PDP with a hint R_1 for the explicit settings as follows: Construct R'_2 and check that (R_1, R'_2) meet the TDP/PDP conditions. If R'_2 meets these conditions then (R_1, R'_2) is a solution for TDP/PDP. Otherwise, there is no solution with R_1 as a hint. From this observation we get:

Theorem 6.3. *TDP/PDP with a hint R_1 for explicit relations are in PTIME.*

The definition of R'_2 can also solve the symbolic setting with a hint. Thus the following result allows the use of QBF solvers for finding solution for R_1 .

Theorem 6.4. *TDP/PDP with a hint C_{R_1} for symbolic relations are in Π_3^P .*

Proof. For the symbolic setting, when R and R_1 are given as circuits C_R and C_{R_1} , resp., we can use the definitions of $C_{R'_2}$ to verify the TDP/PDP conditions: for TDP verify the following formulas (1) $\forall \vec{b}, \vec{i} \exists \vec{o} (C_{R_1}(\vec{i}, \vec{b}) \rightarrow C_{R'_2}(\vec{b}, \vec{o}))$, and (2) $\forall \vec{i}, \vec{o} (\exists \vec{b} (C_{R_1}(\vec{i}, \vec{b}) \wedge C_{R'_2}(\vec{b}, \vec{o})) \leftrightarrow C_R(\vec{i}, \vec{o}))$. For PDP in addition to Formula (1) we also need to verify (3) $\forall \vec{i} (\exists \vec{b}, \vec{o} (C_{R_1}(\vec{i}, \vec{b}) \wedge C_{R'_2}(\vec{b}, \vec{o})) \leftrightarrow \exists \vec{o}' C_R(\vec{i}, \vec{o}'))$, and (4) $\forall \vec{i}, \vec{o} (\exists \vec{b} (C_{R_1}(\vec{i}, \vec{b}) \wedge C_{R'_2}(\vec{b}, \vec{o})) \rightarrow C_R(\vec{i}, \vec{o}))$. In all these, $C_{R'_2}$ need not be constructed but can be described by the formula $(\forall \vec{i}) (C_{R_1}(\vec{i}, \vec{b}) \rightarrow C_R(\vec{i}, \vec{o}))$.¹ $\square$

For the automatic relation setting we have the following result for TDP/PDP with a hint R_1 , where we assume that R_1 is given as a DFA (though users are more likely to use regular expressions).

Theorem 6.5. *TDP/PDP with a hint $\mathcal{A}_{R_1}$ for automatic relations are in EXPTIME.*

Proof. We start by recalling some basic automata constructions. Let A, B, C be (possibly infinite) domains. If $\mathcal{A}$ is an automaton that describes a relation $R \subseteq A \times B$ (in the sense of Section 5.3) then an automaton $Img(\mathcal{A})$ that describes $Img(R)$ is obtained by eliminating the B labels from $\mathcal{A}$ and determinizing. Similarly an automaton $Dom(\mathcal{A})$ that describes $Dom(R)$ is obtained by eliminating the A labels from $\mathcal{A}$ and determinizing. From automaton $\mathcal{A}_1$ that describes a relation $R_1 \subseteq A \times B$, and an automaton $\mathcal{A}_2$ that describes a relation $R_2 \subseteq B \times C$, we obtain an automaton denoted by $\mathcal{A}_1 \circ \mathcal{A}_2$ that describes $R_1 \circ R_2$ by taking the product automaton of $\mathcal{A}_1$ and $\mathcal{A}_2$, eliminate the labels (a, b, b', c) in which $b \neq b'$, then eliminating the B letters from the remaining labels and determinizing. Finally, to check if a DFA $\mathcal{A}$ is contained in a DFA $\mathcal{A}'$ we check for emptiness the automaton $\mathcal{A} \cap \neg \mathcal{A}'$.

Now for the proof, assume that $\mathcal{A}_R = (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{O}}, Q, q^0, \delta, F)$ is a DFA that describes R and $\mathcal{A}_{R_1} = (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}}, Q_1, q_1^0, \delta_1, F_1)$ is a DFA hint. We show how to construct a DFA $\mathcal{A}_2 = (\Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}}, Q_2, q_2^0, \delta_2, F_2)$ that describes the maximum relation R'_2 . We define $\mathcal{A}' = (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{O}}, Q_1 \times Q, (q_1^0, q^0), \delta_1 \times \delta, F')$ as the product automaton of $\mathcal{A}_{R_1}$ and $\mathcal{A}_R$. We re-define the accepting states $F' = F \times F_1$ by setting $F' = \{(q, q') | q \in F_1 \rightarrow q' \in F\}$. For the transition function, we first extract from $\mathcal{A}'$ all the (i, b, i', o) edges in which $(i \neq i')$. Next we delete the alphabets $\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{I}}$ from all the labels in $\mathcal{A}'$. This results in a non-deterministic automaton, that we determinize (with possibly an exponential blow-up) in the standard way through the subset construction, with the one exception that we set every super-state (in the subset construction) to be accepting if and only if *all* of its elements are accepting states. This results in a DFA $\mathcal{A}_2$ over $\Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}}$ that describes the maximum relation R'_2 .

Finally we need to check that $(\mathcal{A}_{R_1}, \mathcal{A}_2)$ meet the TDP/PDP conditions. For that by standard operations we check for TDP that $L(Dom(\mathcal{A}_{R_1})) \subseteq L(Img(\mathcal{A}_2))$, and $L(\mathcal{A}_{R_1} \circ \mathcal{A}_2) = L(\mathcal{A})$. For PDP check that $L(Dom(\mathcal{A}_{R_1})) \subseteq L(Img(\mathcal{A}_2))$, $L(Dom(\mathcal{A}_{R_1} \circ \mathcal{A}_2)) = L(Dom(\mathcal{A}))$ and $L(\mathcal{A}_{R_1} \circ \mathcal{A}_2) \subseteq L(\mathcal{A})$. $\square$

Theorem 6.5 provides an elegant solution for TDP/PDP in the automatic setting, in case one of the component, say, $\mathcal{A}_{R_1}$ is required to be bounded by a given (unary) size k :

¹ By performing quantifier elimination on this formula, $C_{R'_2}$ can be constructed explicitly, perhaps in exponential time. c.f., [16].

guess an automaton $\mathcal{A}_{R_1}$ of size k as a hint R_1 , then use Theorem 6.5. This gives us the following.

Corollary 6.6. *TDP/PDP on automatic relations when one of the component DFAs is bounded by a given (unary) size is in NEXPTIME.*

As another corollary of Theorem 6.5, we can say the following on unary intermediate alphabet.

Corollary 6.7. *TDP/PDP for automatic relations with a unary intermediate alphabet is in EXPTIME.*

Proof. For a unary intermediate alphabet $\Sigma_{\mathcal{B}} = \{b\}$, there is a unique solution R_1 that can be used as a hint. Indeed, since for a solution R_1 we have that $\text{Dom}(R_1) = \text{Dom}(R)$ (see Claim 4.2), then every word of length n in $\text{Dom}(R_1)$ must be paired with the word b^n . Therefore a hint $\mathcal{A}_{R_1}$ is constructed from $\mathcal{A}_R$ by simply replacing every $\Sigma_{\mathcal{O}}^*$ word in $\mathcal{A}_R$ with the unique word of the same length over $\Sigma_{\mathcal{B}}$ (this projection yields an NFA, but determinizing in order to get a DFA as a hint is not necessary, since the proof of Theorem 6.5 can work with $\mathcal{A}_{R_1}$ given as an NFA as well without changing the upper bound). $\square$

Finally, for Strategic PDP in the automatic setting, we assume that the hint is given in the form of a transducer.

Theorem 6.8. *Strategic PDP with a hint T_1 is in EXPTIME.*

Proof. For Strategic PDP with a hint T_1 as a transducer we simplify the construction of Section 5.4. Intuitively we directly construct A_2 by taking the product of A_R and T_1 , then solve non-emptiness for A_2 and obtain a witness transducer T_2 . Since however the resulting A_2 is still an alternating automaton, this construction is still in an exponential time.

In details, we have the following. When T_1 is given as a witness we can construct an automaton $\mathcal{A}'_2$ directly from T_1 and $\mathcal{A}_R^t$. In this case, however, we need to be more careful. What $\mathcal{A}'_2$ does when reading $o \in \Sigma_{\mathcal{O}}$ is for every possible $i \in \Sigma_{\mathcal{I}}$, simulates T_1 to obtain $b \in \Sigma_{\mathcal{B}}$ and then uses $\mathcal{A}_R^t$ to obtain the next states for $\mathcal{A}'_2$ to proceed. Since the letter b can be read by many i 's, this results in an alternating automaton. In details, this is done as follows.

Let $T_1 = (\Sigma_{\mathcal{I}}, \Sigma_{\mathcal{B}}, S, \nu, s_0, \text{Out})$ where $\nu : S \times \Sigma_{\mathcal{I}} \rightarrow S$ is the transition function and $\text{Out} : S \rightarrow \Sigma_{\mathcal{B}}$ is the output function of T_1 . Recall that $\mathcal{A}_R^t = (\Sigma_{\mathcal{I}}, \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}}, Q, q_0, \delta, \alpha)$. Then we construct an automaton $\mathcal{A}'_2 = (\Sigma_{\mathcal{B}}, \Sigma_{\mathcal{O}}, Q \times S, (q_0, s_0), \delta_2, \alpha_2)$ where $\alpha_2 = \alpha \times S$ (i.e. replace every element set $Q' \subseteq Q$ from α with $Q' \times S$), and δ_2 is defined as follows. For $q \in Q, s \in S$ and $o \in \Sigma_{\mathcal{O}}$ we set $\delta_2((q, s), o)$ as follows. For every $i \in \Sigma_{\mathcal{I}}$, let $b = \text{Out}(\nu(s, i))$. Then check $(z_1, \dots, z_{|\Sigma_{\mathcal{I}}|}) \in \delta(q, (b, o))$ and then send $\delta_2((q, s), o)$ to all the $(b, z_i \times \nu(s, i))$. Formally we have:

$$\delta_2((q, s), o) = \bigwedge_{i \in \Sigma_{\mathcal{I}}, b = \text{Out}(\nu(s, i))} \bigvee_{\vec{z} \in \delta(q, (b, o))} (b, z_1 \times \nu(s, i_1)) \wedge \dots \wedge (b, z_{|\Sigma_{\mathcal{I}}|} \times \nu(s, i_{|\Sigma_{\mathcal{I}}|})) \quad (6.1)$$

where $\vec{z} = (z_1, \dots, z_{|\Sigma_{\mathcal{I}}|})$. Then the transducers obtained from $L(\mathcal{A}'_2)$ are exactly those of the form T'_2 such that (T_1, T'_2) meet the Strategic PDP condition.

Note that this construction is also in exponential time due to the emptiness checking of an alternating automaton. Therefore we have that Strategic PDP with hint T_1 is in exponential time. $\square$

6.2. TDP/PDP with hint R_2 . We now discuss the case for TDP/PDP when R_2 is given as a hint. The definition, given below, is similar to that of R_1 being a hint.

Definition 6.9. TDP/PDP problem with a hint R_2 we are given input, output and intermediate domain $\mathcal{I}, \mathcal{O}, \mathcal{B}$, relations $R \subseteq \mathcal{I} \times \mathcal{O}$ and $R_2 \subseteq \mathcal{B} \times \mathcal{O}$. The goal is to find whether there is a relation $R_1 \subseteq \mathcal{I} \times \mathcal{B}$ such that (R_1, R_2) meet the TDP/PDP conditions.

As before, we define the maximum relation R'_1 to be $R'_1 = \{(i, b) | (\exists o(b, o) \in R_2) \wedge (\forall o((b, o) \in R_2 \rightarrow (i, o) \in R))\}$. By an almost identical argument to that of Lemma 6.2 we get the following.

Lemma 6.10. *Every solution for TDP/PDP with a hint R_1 is contained in R'_1 and if there exists such a solution then R'_1 is a solution as well.*

For the explicit setting, construction again is easy. Therefore we have:

Theorem 6.11. *TDP/PDP with a hint R_2 for explicit relations is in PTIME.*

And we also have

Theorem 6.12. *TDP/PDP with a hint C_{R_2} for symbolic relations is in Π_3^P .*

Proof. The arguments for the proof are almost identical to those of Theorem 6.4, with the one change that this time we are given C_{R_2} and we need to verify for TDP the formulas (1) $\forall \vec{b}, \vec{i} \exists \vec{o} (C_{R'_1}(\vec{i}, \vec{b}) \rightarrow C_{R_2}(\vec{b}, \vec{o}))$ and (2) $\forall \vec{i}, \vec{o} \exists \vec{b} ((C_{R'_1}(\vec{i}, \vec{b}) \wedge C_{R_2}(\vec{b}, \vec{o})) \leftrightarrow C_R(\vec{i}, \vec{o}))$. For PDP in addition to Formula (1) we also need to verify (3) $\forall \vec{i} (\exists \vec{b}, \vec{o} (C_{R'_1}(\vec{i}, \vec{b}) \wedge C_{R_2}(\vec{b}, \vec{o})) \leftrightarrow \exists \vec{o} C_R(\vec{i}, \vec{o}))$, and (4) $\forall \vec{i}, \vec{o} (\exists \vec{b} (C_{R'_1}(\vec{i}, \vec{b}) \wedge C_{R_2}(\vec{b}, \vec{o})) \rightarrow C_R(\vec{i}, \vec{o}))$. In all these $C_{R'_1} \equiv (\exists \vec{o} C_{R_2}(\vec{b}, \vec{o}) \wedge \forall \vec{o} C_{R_2}(\vec{b}, \vec{o}) \rightarrow C_R(\vec{i}, \vec{o}))$. $\square$

Next, we can get a similar result to Theorem 6.5 for TDP/PDP for automatic relations with a hint R_2 , but we need to be more careful since R_1 imposes an extra condition of $\text{Img}(R_1) \subseteq \text{Dom}(R_2)$ that is needed to be considered.

Theorem 6.13. *TDP/PDP for automatic relations with a hint $\mathcal{A}_{R_2}$ is in EXPTIME.*

Proof. Assume that $\mathcal{A} = (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{O}}, Q, q^0, \delta, F)$ is a DFA for R , and that $\mathcal{A}_{R_2} = (\Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}}, Q_2, q_2^0, \delta_2, F_2)$ is a DFA for R_2 given as a hint. We show how to construct an automaton $\mathcal{A}_1 = (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}}, Q_1, q_1^0, \delta_1, F_1)$, such that $L(\mathcal{A}_1) = L(R'_1)$ where R'_1 is the maximum relation from Lemma 6.10.

We first construct an automaton that satisfies the second condition from the definition of R'_1 , i.e. for every $\vec{o} \in \Sigma_{\mathcal{O}}$, if $(\vec{b}, \vec{o}) \in R_2$ then $(\vec{i}, \vec{o}) \in R$. Similar to the proof of Theorem 6.5, we define $\mathcal{A}' = (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{O}} \times \Sigma_{\mathcal{B}} \times \Sigma_{\mathcal{O}}, Q \times Q_2, (q^0, q_2^0), \delta \times \delta_2, F')$ to be the product automaton of $\mathcal{A}$ and $\mathcal{A}_{R_2}$. We redefine the accepting states to be $F' = \{(q, q') | q' \in F_2 \rightarrow q \in F\}$. To redefine the transition function, we first extract from $\mathcal{A}'$ all the edges (i, o', b, o) edges in which $(o \neq o')$. Next we delete the alphabets $\Sigma_{\mathcal{O}} \times \Sigma_{\mathcal{O}}$ from all the labels in $\mathcal{A}'$. This results in a non-deterministic automaton, that we determinize in the usual way through the subset construction, with the one exception that we set every super-state (in the subset construction) to be accepting if and only if *all* of its' elements are accepting states. This results in an automaton $\mathcal{A}'_1$ over $\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}}$.

Next, we have to take care of the required condition that if $(\vec{i}, \vec{b})$ is an accepting word in $\mathcal{A}'_1$ then $\vec{b} \in \text{Dom}(R_2)$. For that we first construct $\mathcal{A}'_2 = \text{Dom}(\mathcal{A}_{R_2})$. Then, we define $\mathcal{A}_1$ by taking the product automaton of $\mathcal{A}'_1$ and $\mathcal{A}'_2$ in the sense of $\mathcal{A}_1 = (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{B}} \times$

$\Sigma_{\mathcal{B}}, \mathcal{A}'_1.Q \times \mathcal{A}'_2.Q, (\mathcal{A}'_1.q_0, \mathcal{A}'_2.q_0), \mathcal{A}'_1.\delta \times \mathcal{A}'_2.\delta, \mathcal{A}'_1.F \times \mathcal{A}'_2.F$), extracting all edges labeled (i, b, b') in which $b \neq b'$ and determinizing.

Claim 6.14. $\mathcal{A}_1$ defines the maximum relation R'_1 .

Proof. Let $(\vec{i}, \vec{b})$ be an accepting word in $\mathcal{A}_1$. Then $(\vec{i}, \vec{b}, \vec{b})$ is an accepting word in $\mathcal{A}'_1 \times \mathcal{A}'_2$. Specifically $\vec{b}$ is an accepting word in $\mathcal{A}'_2$, therefore there is a word $\vec{o} \in \Sigma_{\mathcal{O}}$ such that $(\vec{b}, \vec{o})$ is accepting in $\mathcal{A}_{R_2}$ so $\vec{b} \in \text{Dom}(R_2)$. In addition since $(\vec{i}, \vec{b})$ is an accepting word in $\mathcal{A}'_1$ then for every $\vec{o} \in \Sigma_{\mathcal{O}}$ such that $(\vec{b}, \vec{o})$ is in $L(\mathcal{A}_{R_2})$ we have that $(\vec{i}, \vec{o})$ is in $L(\mathcal{A})$. Otherwise we have that $(\vec{b}, \vec{o})$ reaches an accepting state in $\mathcal{A}_{R_2}$, and a non-accepting state in $\mathcal{A}$. Therefore by our construction the word $(\vec{i}, \vec{o}, \vec{b}, \vec{o})$ is not accepting in $\mathcal{A}'$ before removing the $\Sigma_{\mathcal{O}}$ letters, therefore $(\vec{i}, \vec{b})$ is not in $\mathcal{A}'$, a contradiction.

For the other side, assume $(\vec{i}, \vec{b}) \in R'_1$. Then $\vec{b} \in \text{Dom}(R_2)$, therefore $\vec{b} \in L(\mathcal{A}'_2)$. In addition, for every $\vec{o}$ such that $(\vec{b}, \vec{o}) \in R_2$, we have that $(\vec{i}, \vec{o}) \in R$. So if $(\vec{b}, \vec{o}) \in L(\mathcal{A}_{R_2})$ then $(\vec{i}, \vec{o}) \in L(\mathcal{A})$. Therefore for every $\vec{o} \in \Sigma_{\mathcal{O}}$, we have $(\vec{i}, \vec{o}, \vec{b}, \vec{o})$ is accepting in the construction process of $\mathcal{A}'$ before removing the $\Sigma_{\mathcal{O}}$ letters and determinizing. This means that $(\vec{i}, \vec{b}) \in L(\mathcal{A}')$. Therefore all in all we have that $(\vec{i}, \vec{b})$ is accepting in $\mathcal{A}'$ and in $\mathcal{A}'_2$, thus $(\vec{i}, \vec{b})$ is accepting in $\mathcal{A}_1$ as required. $\square$

All is left is to we check that $(\mathcal{A}_1, \mathcal{A}_{R_2})$ meet the condition for TDP/PDP in the same way as in the proof of Theorem 6.5. $\square$

Finally for strategic PDP with a hint T_2 note that when T_2 is given as a hint we do not need to construct $\mathcal{A}_2$, but simply construct $\mathcal{A}_1$ from $\mathcal{A}'^t_R, T_2$ in the exact construction as above. Then we check for emptiness. As before, if $\mathcal{A}_1$ is empty then T_2 is not a good witness. Otherwise, we get a transducer T_1 such that (T_1, T_2) meet the Strategic PDP conditions. Note that this construction yields T_1 in polynomial time therefore we overall have.

Theorem 6.15. *Strategic PDP with a hint T_2 is in PTIME.*

7. DISCUSSION

We studied here a formal model of sequential decomposition, a fundamental concept in computational thinking. We showed that while decomposition is viewed as an approach to tame design complexity, complexity is not so easily tamed and decomposition can be quite difficult when viewed as a computational problem. Human intuition, used to offer hints to the decomposition algorithm, is therefore necessary to tame the complexity of the decomposition problem. The complexity of TDP/PDP in the automatic-relation setting, conjectured to be undecidable, is still open. It is related to other decision problems for automatic relations, cf. [29], and is a subject of future work.

Another line of future work is to explore the orthogonal problem of *parallel* decomposition, in which a task is decomposed into two tasks executed in parallel. This subject has been studied in [1], but the specification of the system there is given by the conjunction of its components' specifications, while we study specifications in terms of overall system input/output relations.

Finally a more practical approach to sequential decomposition can be found in synthesis for symbolic settings. Indeed we showed in Section 5.2 that symbolic TDP/PDP is NEXPTIME-complete, but this stands only when the size of the intermediate domain $\mathcal{B}$ is

given. For the synthesis approach we relax this requirement so we get to choose $\mathcal{B}$. On one hand the problem of solving TDP/PDP becomes uninteresting. On the other, however, The challenge is to find a TDP/PDP solution (R_1, R_2) that allows good synthesis for a given relation R in the sense that the composition of a separate synthesis to R_1 and R_2 can be a synthesis solution to R . A recent work inspired by this approach can be found in [9].

REFERENCES

- [1] M. Abadi and L. Lamport. Decomposing specifications of concurrent systems. In *Proc. IFIP Working Conference on Programming Concepts, Methods and Calculi*, pages 327–340, 1994.
- [2] S. Akshay, T. Antonopoulos, J. Ouaknine, and J. Worrell. Reachability problems for Markov chains. *Inf. Process. Lett.*, 115(2):155–158, 2015.
- [3] N. Amálio and C. Glodt. A tool for visual and formal modelling of software designs. *Sci. Comput. Program.*, 98:52–79, 2015.
- [4] J. L. Balcázar, A. Lozano, and J. Torán. The complexity of algorithmic problems on succinct instances. In *Computer Science*, pages 351–377. Springer, 1992.
- [5] A. Basu, S. Bensalem, M. Bozga, J. Combaz, M. Jaber, T. Nguyen, and J. Sifakis. Rigorous component-based system design using the BIP framework. *IEEE Software*, 28(3):41–48, 2011.
- [6] A. Benveniste, B. Caillaud, A. Ferrari, L. Mangeruca, R. Passerone, and C. Sofronis. Multiple viewpoint contract-based specification and design. In *Proc. 6th International Symposium in Formal Methods for Components and Objects FMCO*, pages 200–225, 2007.
- [7] S. Bliudze, J. Sifakis, M. Bozga, and M. Jaber. Architecture internalisation in BIP. In *Proc. 17th International Symposium on Component-Based Software Engineering*, pages 169–178, 2014.
- [8] A. Bouajjani, B. Jonsson, M. Nilsson, and T. Touili. Regular model checking. In *Int’l Conf. on Computer-Aided Verification*, pages 403–418. Springer, 2000.
- [9] Supratik Chakraborty, Dror Fried, Lucas M. Tabajara, and Moshe Y. Vardi. Functional synthesis via input-output separation. In *2018 Formal Methods in Computer Aided Design, FMCAD 2018*.
- [10] C. Cheng, S. Bensalem, Y. Chen, R. Yan, B. Jobstmann, H. Ruess, C. Buckl, and A. Knoll. Algorithms for synthesizing priorities in component-based systems. In *Proc. 9th International Symposium in Automated Technology for Verification and Analysis ATVA*.
- [11] B. Das, P. Scharpfenecker, and J. Torán. CNF and DNF succinct graph encodings. *Information and Computation*, 2016.
- [12] L. de Alfaro and T. A. Henzinger. Interface theories for component-based design. In *Proc. First International Workshop on Embedded Software EMSOFT*, pages 148–165, 2001.
- [13] U. Fahrenberg and A. Legay. General quantitative specification theories with modal transition systems. *Acta Inf.*, 51(5):261–295, 2014.
- [14] J. Feigenbaum, S. Kannan, M. Y. Vardi, and M. Viswanathan. The complexity of problems on graphs represented as obdds. *Chicago J. Theor. Comput. Sci.*, 1999, 1999.
- [15] L. Fix, N. Francez, and O. Grumberg. Program composition and modular verification. In *Proc. 18th Int. Colloq. on Automata, Languages, and Programming*, pages 93–114, 1991.
- [16] D. Fried, L.M. Tabajara, and M.Y. Vardi. BDD-based Boolean functional synthesis. In *Proc. 28th Int’l Conf. on Computer Aided Verification*, pages 402–421, 2016.
- [17] H. Galperin and A. Wigderson. Succinct representations of graphs. *Information and Control*, 56(3):183–198, 1983.
- [18] M. R. Garey and D. S. Johnson. *Computers and intractability: A guide to the theory of NP-Completeness*. W. H. Freeman, 1979.
- [19] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, 1969.
- [20] J. E. Hopcroft, R. Motwani, and J. D. Ullman. *Introduction to automata theory, languages, and computation - international edition (2. ed)*. Addison-Wesley, 2003.
- [21] O. James. Contentment in graph theory: covering graphs with cliques. In *Indagationes Mathematicae (Proceedings)*, volume 80, pages 406–424. Elsevier, 1977.
- [22] D.S. Johnson. The NP-completeness column: An ongoing guide. *Journal of algorithms*, 8(2):285–303, 1987.

- [23] R. M. Karp. Reducibility among combinatorial problems. In *Proc. of a symposium on the Complexity of Computer Computations*, pages 85–103, 1972.
- [24] B. Khoussainov and A. Nerode. Automatic presentations of structures. In *Proc. Workshop on Logic and computational complexity*, LNCS 960, pages 367–392. Springer, 1995.
- [25] D. Kozen. *Automata and computability*. Undergraduate texts in computer science. Springer, 1997.
- [26] O. Kupferman and M.Y. Vardi. Synthesis with incomplete informatio. *Advances in Temporal Logic*, 16:109–127, 2000.
- [27] O. Kupferman and M.Y. Vardi. Synthesizing distributed systems. In *Proc. 16th IEEE Symp. on Logic in Computer Science*, pages 389–398, 2001.
- [28] E. Kushilevitz and N. Nisan. *Communication complexity*. Cambridge University Press, 1997.
- [29] D. Kuske and M. Lohrey. Some natural decision problems in automatic graphs. *J. of Symbolic Logic*, pages 678–710, 2010.
- [30] K. G. Larsen and L. Xinxin. Equation solving using modal transition systems. In *Proceedings of the Fifth Annual Symposium on Logic in Computer Science (LICS)*, pages 108–117, 1990.
- [31] K. Guldstrand Larsen. Modal specifications. In *Proc. Automatic Verification Methods for Finite State Systems, International Workshop*, pages 232–246, 1989.
- [32] Y. Lustig and M.Y. Vardi. Synthesis from component libraries. *Software Tools for Technology Transfe*, 15(5-6):603–618, 2013.
- [33] N. Lynch and M. R. Tuttle. An introduction to Input/Output automata. *CWI-quarterly*, 2(3), 1989.
- [34] J. Ouaknine and J. Worrell. On the Positivity problem for simple linear recurrence sequences. In *Proc. ICALP*, volume 8573 of *LNCS*. Springer, 2014.
- [35] J. Ouaknine and J. Worrell. On linear recurrence sequences and loop termination. *SIGLOG News*, 2(2):4–13, 2015.
- [36] C. H. Papadimitriou and M. Yannakakis. A note on succinct representations of graphs. *Information and Control*, 71(3):181–185, 1986.
- [37] D.L. Parnas. On the criteria to be used in decomposing systems into modules. *Commun. of the ACM*, 15(12):1053–1058, 1972.
- [38] A. Pnueli. The temporal logic of programs. In *Proc. 18th Annual Symposium on Foundations of Computer Science*, pages 46–57, 1977.
- [39] A. Pnueli and R. Rosner. Distributed reactive systems are hard to synthesize. In *Proc. 31st Annual Symposium on Foundations of Computer Science*, pages 746–757, 1990.
- [40] K. Rath, V. Choppella, and S.D. Johnson. Decomposition of sequential behavior using interface specification and complementation. *VLSI Design*, 3(3-4):347–358, 1995.
- [41] S. Rubin. *Automatic structures*. PhD thesis, University of Auckland, New Zealand, 2004.
- [42] M. Soittola. On DOL synthesis problem. In *Automata, Languages, Development*. North-Holland, 1976.
- [43] H. Veith. Languages represented by boolean formulas. *Inf. Process. Lett.*, 63(5):251–256, 1997.
- [44] J. M. Wing. Computational thinking. *Commun. ACM*, 49(3):33–35, 2006.
- [45] J. M. Wing. Computational thinking. In *Proc. IEEE Symposium on Visual Languages and Human-Centric Computing*, page 3, 2011.