

Automatic Verification of Data-Centric Business Processes

Alin Deutsch
UC San Diego

Richard Hull
IBM Yorktown Heights

Fabio Patrizi
Sapienza Univ. of Rome

Victor Vianu
UC San Diego

Abstract

We formalize and study business process systems that are centered around "business artifacts", or simply "artifacts". This approach focuses on data records, known as artifacts, that correspond to key business-relevant objects, and that flow through a business process specified by a set of services. The artifact-centric approach has been introduced by IBM, and has enabled significant improvements to the operations of medium- and large-sized businesses.

In this paper, artifacts carry attribute records and internal state relations, that services can consult and update. In addition, services can access an underlying database and can introduce new values from an infinite domain, thus modeling external inputs or partially specified processes described by pre-and-post conditions. The services are associated to the artifacts using declarative, condition-action style rules.

We consider the problem of statically verifying whether all runs of an artifact system satisfy desirable correctness properties expressed in a first-order extension of linear-time temporal logic. We map the boundaries of decidability for the verification problem and provide its complexity. The technical challenge to static verification lies in the fact that artifact systems are infinite-state systems, as the domain of the data is infinite. We identify an expressive class of artifact systems for which verification is nonetheless decidable. Remarkably, the complexity of verification is PSPACE-complete, which is no worse than classical finite-state model checking.

This investigation builds upon previous work on verification of data-driven Web services and ASM transducers, while addressing significant new technical challenges raised by the artifact model.

1 Introduction

Businesses and other organizations increasingly rely on business process management, and in particular the management of electronic workflows underlying business processes. While most workflow is still organized around relatively flat process-centric models, over the past several years a *data-centric* approach to workflow has emerged. A watershed paper in this area is [28], which introduces the *artifact-centric* approach to workflow modeling. This approach focuses on data records, known as "business artifacts" or simply "artifacts", that correspond to key business-relevant objects, their life cycles, and how/when services (a.k.a. tasks) are invoked on the artifacts. This approach provides a simple and robust structure for workflow, and has been demonstrated in practice to yield substantial improvements in the implementation of business transformations [4].

From the formal perspective, little is understood about artifact-centric (and other data-centric) workflow. Citations [15, 16, 6] provide preliminary investigations into the analysis of artifact-centric workflow in restricted settings. The current paper develops static analysis techniques in the context of substantially richer declarative artifact-centric workflows, in which each artifact carries, in addition to the data record containing business-relevant values, a full relational state used in the internal control of the business process. The focus is on decision problems for a first-order extension of linear-time temporal logic. In general such decision problems are undecidable, but the paper identifies a large, useful class of workflows for which the complexity of verification is PSPACE-complete, which is no worse than classical finite-state model checking.

Following [7], an artifact-centric workflow model typically focuses on (a) the business artifacts, (b) the (macro-)life cycle of the artifacts, (c) the services (a.k.a. tasks) that operate on the artifacts, and (d) the mechanisms whereby services are associated to the artifacts. In this paper each artifact type will involve a family of attributes along with one relational state (a simple generalization permits multiple relational states). Intuitively, the artifact starts with just a few of the attributes defined (initialized), and the invoked services fill in, or overwrite, the artifact's attributes. In practice, the primary operations of a business may involve tens of artifacts for small-to-medium-size businesses, and hundreds of artifacts for large businesses. The macro-lifecycle of the artifacts is intended to capture the key business-relevant *stages* (or states) in the lifecycle of an artifact. Stage transitions are typically specified using a finite-state machine, often with a handful to tens of states.

The current paper specifies services and their association to artifacts in a declarative manner, using input parameters, output parameters, pre-conditions, and post-conditions. This model is inspired by the model of [6], and more broadly by the field of semantic web services [25] (where post-conditions are called "conditional effects"). The use of post-conditions permits non-determinism in the outcome of a service, as is typically the case, for example, in a business process service in which a human makes a final determination about the value of an attribute while satisfying certain constraints. Also following [6], in this paper the movement of artifacts from one stage to another is specified declaratively, in our case via state update rules (known as condition-action rules in [6]). The model of [6] permits only attributes, and the analysis focuses only on whether these attributes are defined or undefined (so their value is abstracted away). In the current paper we handle both attributes and relational states. The service and property specifications manipulate the data values they hold, and can compare them according to a dense linear order. Further, we include here a static database which can be accessed (but not updated) during the processing of artifacts by the services' state update rules and pre- and post-conditions.

The workflow model used in this paper is illustrated with a running example that models a scenario where a manufacturer fills customer purchase orders, negotiating the price of each item on a case-by-case basis. The example is introduced in Example 2.4 and presented in full in Appendix A.

This paper considers the problem of statically verifying whether all runs of an artifact system satisfy desirable correctness properties expressed in a first-order extension of linear-time temporal logic called LTL-FO . This language can express a wide variety of properties pertaining to the consistency of the specification (e.g two Boolean flags are mutually exclusive at every step of the business process), or to the policies implemented by a business process. For instance, in the running example, one wishes to guarantee the following:

If the customer's status is not *preferred* and the credit rating is worse than *good*, then before accepting an order for a product with final negotiated price lower than the manufacturer's desired price, explicit approval from a human executive must be requested.

We also show how other common analysis tasks for business processes can be reduced to verification of LTL-FO properties.

The main technical challenge to static verification lies in the fact that the artifact systems studied here are infinite-state systems, as the domain of the data is infinite. In the general case, testing correctness of properties is undecidable. We map the boundaries of decidability for the verification problem, and identify a class of restrictions such that (a) artifact systems and properties that lie within this class are decidable, and (b) relaxing any of the restrictions leads to undecidability. For the restricted setting, the decision problem is PSPACE-complete. As will be seen, the running example obeys the restrictions, thus illustrating that they are not prohibitive for practical scenarios.

Further related work As mentioned above, artifacts and related notions have been discussed in the research literature for several years now. The specific notion of artifact, along with specification of key stages in its life-cycle, was first introduced in [28] and was further studied, from both practical and theoretical perspectives, in [4, 5, 15, 16, 6, 23, 20, 22]. Some key roots of the artifact-centric approach are present in adaptive objects

[21], adaptive business objects [26], business entities, and “document-driven” workflow [34]. The notion of documents as in document engineering [17] is focused on certain aspects of artifacts, namely the artifact data itself and how it can be used to facilitate communication between sub-organizations in the course of workflow processing. The Vortex workflow framework [19, 12, 18] is also data centric, and provides a declarative framework for specifying if and when workflow services are to be applied to a given artifact. More recently, [2] has studied automatic verification in the context of workflow based on Active XML documents.

Work on formal analysis of artifact-centric business processes in restricted contexts has been reported in [15, 16, 6]. Properties investigated in these studies include reachability [15, 16], general temporal constraints [16], and the existence of complete execution or dead end [6]. Citations [15, 16] are focused on an essentially procedural version of artifact-centric workflow, and [6] is the first to study a declarative version. For the variants considered in each paper, verification is generally undecidable; decidability results were obtained when rather severe restrictions are placed, e.g., restricting all guards on state transitions to be “true” [15], restricting to bounded domains [16, 6], or restricting the language for conditions to refer only to artifacts (and not their attribute values) [16]. None of the above papers permit an arbitrary external database, separate from the artifacts, in their frameworks.

The OWL-S proposal [25, 24] describes the semantics of services with input, output, pre-condition, and post-conditions (known there as *conditional effects*). In that work, the pre-conditions and effects refer to *fluents*, that is predicates whose values can change over time. These are used to model evolving databases, for instance for flight reservations, bank accounts, and warehouse inventories. The declarative artifact-centric approach to workflow modeling used here is closely related to that of semantic web services in general, and OWL-S in particular.

Static analysis for semantic web services is considered in [27], but in a context restricted to finite domains.

The work [10] studies static verification of data-driven Web services that interact with external users through a Web browser interface and generate Web pages dynamically by queries on an underlying database. The study identifies decidable cases of the problem of verifying if all runs of a Web service satisfy a correctness property specified as a sentence in LTL-FO, the language of first order logic extended with linear-time temporal logic operators, which we adopt also in this paper. Similar extensions have been previously used in various contexts [13, 1, 32, 10, 11]. The model studied in [10] extends prior formalisms for specifying electronic commerce applications with additional features that turn out to be essential for describing Web applications. Its immediate ancestor is the ASM transducer [32, 31], a more remote one is the relational transducer [3]. The artifact system model could conceptually (if not naturally) be encoded into the extended ASM transducer model of [10]. However, this would not yield a proof of the results in this paper, because business process modeling requires two non-trivial extensions. First, runs of artifact systems must be allowed to use infinitely many domain values in order to model arbitrary inputs from external users or partially specified processes described by pre- and post-conditions (unlike transducers, where the domain of each run is restricted to the active domain of the finite database). Second, the underlying domain is ordered, which turns out to be a key feature in writing practically useful pre- and post-conditions. These extensions render the proof of decidability of verification considerably more involved.

Paper outline Our model of artifact systems and the language LTL-FO are introduced in Section 2, together with our running example. Section 3 states the restrictions needed for decidability of verification, and provides the main decidability result. In Section 4, the restrictions are shown to be tight by considering several relaxations that lead to undecidability of verification. Applications of the main verification results to other business process analysis tasks are provided in Section 5. We end with brief conclusions. An Appendix contains the full running example and most proofs, including the proof of the main decidability result.

2 Framework

We introduce here our model and basic definitions and notation.

We assume fixed an infinite, countable domain D equipped with a total dense order \leq with no endpoints. As usual, a database schema \mathcal{D} consists of a finite set of relation symbols with specified arities. The arity of relation R is denoted $a(R)$. An instance, or interpretation, over a database schema, is a mapping associating to each relation symbol R of the schema a finite relation over D , of arity $a(R)$. We assume familiarity with First-Order logic (FO) over database schemas. Given a schema \mathcal{D} , $\mathcal{L}_{\mathcal{D}}$ denotes the set of FO formulas over $\mathcal{D} \cup \{=, \leq\}$ ($=$ and \leq are built-in relations over D). If $\varphi(\bar{x})$ is an FO formula with free variables \bar{x} , and \bar{u} is a tuple over D of the same arity as \bar{x} , we denote by $\varphi(\bar{x} \leftarrow \bar{u})$ the sentence obtained by substituting \bar{u} for \bar{x} in $\varphi(\bar{x})$. When there is no ambiguity, we sometimes denote $\varphi(\bar{x} \leftarrow \bar{u})$ simply by $\varphi(\bar{u})$. Note that, since D is infinite, an FO formula $\varphi(\bar{x})$ may be satisfied by infinitely many tuples \bar{u} over D (so may define an infinite relation). Finiteness and effective evaluation can be guaranteed by using the *active domain semantics*, in which the domain is restricted to the set of elements occurring in the given instance (sometimes augmented with a specified finite set of constants in D , by default empty). For an instance I , we denote its active domain by $adom(I)$. We assume unrestricted semantics unless otherwise specified.

The artifact model uses a specific notion of class, schema and instance, defined next.

Definition 2.1 An *artifact class* is a pair $\mathcal{C} = \langle R, S \rangle$ where R and S are two relation symbols. An *instance* of \mathcal{C} is a pair $C = \langle R, S \rangle$, where (i) R , called *attribute relation*, is an interpretation of R containing exactly one tuple over D , and (ii) S , called *state relation*, is a finite interpretation of S over D .

We also refer to an *artifact instance of class \mathcal{C}* as *artifact instance*, or simply *artifact* when the class is clear from the context or irrelevant.

Definition 2.2 An *artifact schema* is a tuple $\mathcal{A} = \langle \mathcal{C}_1, \dots, \mathcal{C}_n, \mathcal{DB} \rangle$ where each $\mathcal{C}_i = \langle R_i, S_i \rangle$ is an artifact class, \mathcal{DB} is a relational schema, and $\mathcal{C}_i, \mathcal{C}_j$, and \mathcal{DB} have no relation symbols in common for $i \neq j$.

By slight abuse, we sometimes identify an artifact schema \mathcal{A} as above with the relational schema $\mathcal{DB}_{\mathcal{A}} = \mathcal{DB} \cup \{R_i, S_i \mid 1 \leq i \leq n\}$.

An instance of an artifact schema is a tuple of class instances, each corresponding to an artifact class, plus a database instance:

Definition 2.3 An *instance* of an artifact schema $\mathcal{A} = \langle \mathcal{C}_1, \dots, \mathcal{C}_n, \mathcal{DB} \rangle$ is a tuple $A = \langle C_1, \dots, C_n, DB \rangle$, where C_i is an instance of \mathcal{C}_i and DB is an instance of \mathcal{DB} over D .

Again by slight abuse, we identify each instance $A = \langle C_1, \dots, C_n, DB \rangle$ of \mathcal{A} with the relational instance $DB \cup \{R_i, S_i \mid 1 \leq i \leq n\}$ over schema $\mathcal{DB}_{\mathcal{A}}$. Let \mathcal{A} be an artifact schema and $\mathcal{DB}_{\mathcal{A}}$ its relational schema. Given an artifact instance over \mathcal{A} , the semantics of formulas in $\mathcal{L}_{\mathcal{A}}$ is the standard semantics on the associated relational instance over $\mathcal{DB}_{\mathcal{A}}$.

Example 2.4 We illustrate the expressive power of the artifact model by specifying a scenario where a manufacturer fills customer purchase orders, negotiating the price of each line item on a case-by-case basis. We focus on two artifacts manipulated by the negotiation process, ORDER and QUOTE.

During the workflow, the customer repeatedly adds new line items into (or updates existing ones in) the purchase order modeled by the ORDER artifact. Each line item specifies a product and its quantity. Every tentative line item spawns a negotiation process, in which manufacturer and customer complete rounds of declaring ask and bid prices, until agreement is reached or the negotiation fails. The prices at every round are stored in the QUOTE artifact, which also holds the manufacturer's initially desired price, the lowest bid he is willing to

entertain, and the final negotiated price. Once the negotiation on a tentative line item succeeds, its outcome is scrutinized by a human executive working for the manufacturer. Upon the executive’s approval, the line item is included into the purchase order. During the negotiation, the manufacturer consults an underlying database, which lists information about available products (e.g. manufacturing cost) and about customers (e.g. credit rating and status).

The corresponding artifact system $\Gamma_{ex} = \langle \mathcal{A}, \Sigma \rangle$ is partially described here and in Example 2.9 (see Appendix for the full specification).

The artifact schema is $\mathcal{A} = \langle ORDER, QUOTE, \mathcal{DB} \rangle$, detailed as follows.

$\mathcal{DB} = \langle \text{PRODUCT}, \text{CUSTOMER} \rangle$ is the database schema, where:

- $\text{PRODUCT}(prod_id, manufacturing_cost, min_order_qty)$ lists product manufacturing cost and minimum order quantity, and
- $\text{CUSTOMER}(customer_id, status, credit_rating)$ lists customer status and credit rating.

$ORDER = \langle R_O, line_items, in_process, done \rangle$

is the artifact class containing the information about a customer’s order.

- $R_O(order\#, customer_id, need_by, li_prod, li_qty)$ is the attribute relation holding the order number, the identifier of the customer who placed the order, the day it is needed by. The role of attributes li_prod and li_qty is described later.
- $line_items(prod_id, qty)$ is a state relation that acts as a “shopping cart” holding the collection of line items requested so far.
- $in_process$ and $done$ are nullary state relations (Boolean flags) keeping track of the stage the artifact is in.¹

The intention is that, in stage $in_process$, the customer repeatedly updates the shopping cart by filling an individual, tentative line item into attributes li_prod and li_qty . Subsequently, this line item is inserted into $line_items$ provided the price negotiation succeeds. When the customer completes the purchase order, the $ORDER$ artifact transitions to stage $done$.

$QUOTE = \langle R_Q, li_quotes, idle, desired_price_calc, negotiation, approval_pending, archive \rangle$

is the artifact class modeling quotes, with:

- $R_Q(order\#, desired_price, lowest_acceptable_price, ask, bid, final_price, approved, li_prod, li_qty, manufacturing_cost)$ is the attribute relation.
- $li_quotes(prod_id, qty, price)$ is a state relation storing the line item with the final negotiated price quotes.
- $idle, desired_price_calc, negotiation, approval_pending,$ and $archive$ are nullary state relations keeping track of the stage the artifact is in.

¹Artifact class $ORDER$ illustrates an extension of Definition 2.1 that allows several state relations. This extension is for convenience only: it is easy to show a reduction from multiple-state artifacts to single-state artifacts that preserves our decidability result. The proof is similar to the proof of Lemma B.1.

When inactive, the QUOTE artifact is in state `idle`, but moves to `desired_price_calc` as soon as the customer fills in the product id and quantity of a line item. In this stage, `desired_price` attribute is set (from the manufacturer's point of view), possibly taking into account the `need_by` date attribute in the corresponding ORDER artifact and the manufacturing cost listed in the PRODUCT database. During the ensuing `negotiation` stage, the ask and bid prices are repeatedly set (in attribute `ask` by the manufacturer, respectively `bid` by the customer) until a final price is established and recorded in attribute `final_price`, or the negotiation fails. Final prices may require approval by a human executive who works for the manufacturer. While approval is awaited, the QUOTE artifact is in stage `approval_pending`. Approval is granted by setting Boolean attribute `approved`. Approved final prices are then archived in state relation `li_quotes` (while the QUOTE artifact is in stage `archive`). \square

We now define the syntax of services. It will be useful to associate to each attribute relation R of an artifact schema \mathcal{A} a fixed sequence \bar{x}_R of distinct variables of length $a(R)$.

Definition 2.5 A service σ over an artifact schema \mathcal{A} is a tuple $\sigma = \langle \pi, \psi, \mathcal{S} \rangle$ where:

- π , called *pre-condition*, is a sentence in $\mathcal{L}_{\mathcal{A}}$;
- ψ , called *post-condition*, is a formula in $\mathcal{L}_{\mathcal{A}}$, with free variables $\{\bar{x}_R \mid R \text{ is an attribute relation of a class in } \mathcal{A}\}$;
- \mathcal{S} is a set of *state rules* containing, for each state relation S of \mathcal{A} , one, both or none of the following rules:
 - $S(\bar{x}) \leftarrow \phi_S^+(\bar{x})$;
 - $\neg S(\bar{x}) \leftarrow \phi_S^-(\bar{x})$;

where $\phi_S^+(\bar{x})$ and $\phi_S^-(\bar{x})$ are $\mathcal{L}_{\mathcal{A}}$ -formulas with free variables \bar{x} s.t. $|\bar{x}| = a(S)$.

Definition 2.6 An artifact system is a pair $\Gamma = \langle \mathcal{A}, \Sigma \rangle$, where \mathcal{A} is an artifact schema and Σ is a non-empty set of services over \mathcal{A} .

We next define the semantics of services. We begin with the notion of possible successor of a given artifact instance with respect to a service.

Definition 2.7 Let $\sigma = \langle \pi, \psi, \mathcal{S} \rangle$ be a service over artifact schema \mathcal{A} . Let A and A' be instances of \mathcal{A} . We say that A' is a *possible successor* of A with respect to σ (denoted $A \xrightarrow{\sigma} A'$) if the following hold:

1. $A \models \pi$;
2. $A'|\mathcal{DB} = A|\mathcal{DB}$;
3. if \bar{u}_R is the content of the attribute relation R of \mathcal{A} in A' , then A satisfies the post-condition ψ where each \bar{x}_R is replaced by \bar{u}_R ;
4. for each state relation S of \mathcal{A} and tuple \bar{u} over $adom(A)$ of arity $a(S)$, $A' \models S(\bar{u})$ iff

$$A \models (\phi_S^+(\bar{u}) \wedge \neg \phi_S^-(\bar{u})) \vee (S(\bar{u}) \wedge \phi_S^+(\bar{u}) \wedge \phi_S^-(\bar{u})) \vee (S(\bar{u}) \wedge \neg \phi_S^+(\bar{u}) \wedge \neg \phi_S^-(\bar{u}))$$

where $\phi_S^+(\bar{u})$ and $\phi_S^-(\bar{u})$ are interpreted under active domain semantics, and are taken to be false if the respective rule is not provided.

Note that, according to (2) in Definition 2.7, services do not update the database contents. Instead, the updatable data is carried by the artifacts themselves, as attribute and state relations. This distinction between the static and updatable portions of the data is convenient for technical reasons, as it is used in formulating the restrictions needed for verification (see Section 3). Note that, if so desired, one can make the entire database updatable by turning it into a state. Also observe that the distinction between state and database is only conceptual, and does not preclude implementing all relations within the same DBMS.

We next define the notion of run of an artifact system $\Gamma = \langle \mathcal{A}, \Sigma \rangle$. An *initial instance* of Γ is an artifact instance over \mathcal{A} whose states are empty.

Definition 2.8 A *run* of an artifact system $\Gamma = \langle \mathcal{A}, \Sigma \rangle$ is an infinite sequence $\rho = \{\rho_i\}_{i \geq 0}$ of artifact instances over \mathcal{A} (also called *configurations*) such that:

- ρ_0 is an initial instance of Γ ;
- for each $i \geq 0$, $\rho_i \xrightarrow{\sigma} \rho_{i+1}$ for some $\sigma \in \Sigma$.

A *pre-run* is a finite sequence $\{\rho_i\}_{0 \leq i \leq n}$ satisfying the same conditions as above for $i < n$. We say that a pre-run is *blocking* if its last configuration has no possible successor.

Example 2.9 Continuing Example 2.4, we show how to model the operations allowed on artifacts by the set Σ of available services. Due to space constraints, we relegate most of the specification of Σ to Appendix A, focusing here on the service that models the negotiation process. To illustrate the artifact model's natural ability to specify processes at different levels of abstraction, we describe the negotiation process at two levels. In a first, coarser cut, the process is abstracted as service *abstract_negotiation* = $\langle \pi^{an}, \psi^{an}, \mathcal{S}^{an} \rangle$ about which we only know that the final price is reached when the ask and bid prices coincide, and that it is guaranteed to lie between the allowed margins stored in attributes *desired_price* and *lowest_acceptable_price* of artifact QUOTE. The specification of this service is relatively simple and given in Appendix A. Alternatively, we show below service *refined_negotiation* = $\langle \pi^{rn}, \psi^{rn}, \mathcal{S}^{rn} \rangle$ which refines the negotiation process all the way to the level of individual negotiation rounds, each of which sets the current ask and bid prices.

Conventions We adopt the following conventions:

- (i) We model uninitialized attributes by setting them to the reserved constant ω .
- (ii) We model Boolean states by nullary state relations, and drop the parentheses from atoms using them: $\mathbf{S}()$ becomes \mathbf{S} . We assume the usual encoding of *true* as the singleton nullary relation, and *false* as the empty nullary relation. In particular, all Boolean states are initially *false* (since all state relations are initially empty).
- (iii) For convenience, we use the following syntactic sugar for post-conditions: we write post-conditions as non-Horn rules $h(\bar{x}) := b(\bar{y})$ where the *head* h is a conjunction of atoms over attribute relations in \mathcal{A} , with variables \bar{x} , and the *body* b is a formula in $\mathcal{L}_{\mathcal{A}}$ with free variables \bar{y} , where $\bar{y} \subseteq \bar{x}$. The semantics is that whenever $A \xrightarrow{\sigma} A'$ holds, $A' \models h(\bar{x} \leftarrow \bar{u})$ for some tuple \bar{u} , and $A \models b(\bar{y} \leftarrow \bar{u}|\bar{y})$. Moreover, artifact relations not mentioned in h remain unchanged. Clearly, this syntactic sugar can be simulated by the official post-conditions, and conversely.

We now describe service *refined_negotiation* = $\langle \pi^{rn}, \psi^{rn}, \mathcal{S}^{rn} \rangle$:

The pre-condition

$$\pi^{rn} \doteq \text{negotiation}$$

ensures that the service applies only as long as the Boolean state flag *negotiation* is set in the QUOTE artifact.

Post-condition ψ^{rn} is given as

$$\begin{aligned}
R_Q(o, d, l, a, b, f, app, p, q, m) := & \\
& (\exists a', b' R_Q(o, d, l, a', b', \omega, app, p, q, m) \wedge a' \neq b' \wedge l \leq a \leq a' \wedge b' \leq b \wedge f = \omega) \\
& \vee \\
& (R_Q(o, d, l, a, b, \omega, app, p, q, m) \wedge a = b = f).
\end{aligned}$$

According to the first disjunct, the negotiation is well-formed, i.e. the bid never exceeds the *ask* price, and in each round, asking prices a never increase while bids b never decrease. Moreover, as long as *ask* and *bid* price differ, the final price remains undefined (equal to ω). Notice that the values of a and b are otherwise unconstrained, being simply drawn from the infinite domain. This reflects the fact that they are external input from the manufacturer, respectively customer. The second disjunct states that once *ask* price a and *bid* price b coincide, the final price f is automatically set to the common value.

S^{rn} contains rules that, upon detecting successful negotiation, switch the QUOTE artifact to stage `approval_pending` if the customer does not enjoy preferred status with excellent credit. If he does, then the approval is short-circuited and the QUOTE goes directly to stage `archive`. The negotiation is successful when the *ask* and *bid* prices agree.

$$\begin{aligned}
\text{approval_pending} & \leftarrow \exists o, d, l, a, f, app, p, q, m R_Q(o, d, l, a, a, f, app, q, m) \\
& \quad \wedge \exists c, n R_O(o, c, n, p, q) \wedge \neg \text{CUSTOMER}(c, \text{"preferred"}, \text{"excellent"}) \\
\text{archive} & \leftarrow \exists o, d, l, a, f, app, p, q, m R_Q(o, d, l, a, a, f, app, q, m) \\
& \quad \wedge \exists c, n R_O(o, c, n, p, q) \wedge \text{CUSTOMER}(c, \text{"preferred"}, \text{"excellent"}) \\
\neg \text{negotiation} & \leftarrow \exists o, d, l, a, f, app, p, q, m R_Q(o, d, l, a, a, f, app, p, q, m)
\end{aligned}$$

Note that state flag `negotiation` must be set before the state update rules execute (since pre-condition π^{rn} is satisfied). If neither of the state rule bodies is satisfied, then according to the possible successor semantics, the `negotiation` flag remains set, enabling another negotiation round. \square

One of the points illustrated by Example 2.9 is that the artifact model is particularly well-suited for expressing a wide spectrum of abstraction levels desired in specification. This is shown by the two specifications of the negotiation process, one refining it down to individual rounds, the other abstracting it to an atomic sub-task with a post-condition on its outcome. In practice, the motivation for abstraction ranges from lack of information about an external process provided by an autonomous third party as a black box with pre- and post-execution guarantees, to modeling non-deterministic processes governed by chance or human agents rather than by program. There are also technical reasons, such as the undecidability of verification in the presence of arithmetic (as is the case in many settings, including ours). In all these cases, abstracted sub-processes can be naturally modeled as services, leveraging the non-determinism in their post-conditions.

In order to specify temporal properties of runs, we use an extension of linear-time temporal logic (LTL). Recall that LTL is propositional logic augmented with temporal operators such as **X** (next), **U** (until), **G** (always) and **F** (eventually). Essentially, the extension we use, denoted LTL-FO, is obtained from LTL by replacing propositions by FO statements about individual artifact instances in the run. The different statements may share variables that are universally quantified at the end. Similar extensions have previously been used in various contexts [13, 1, 32, 10, 11].

Definition 2.10 The language LTL-FO (first-order linear-time temporal logic) is obtained by closing FO under negation, disjunction, and the following formula formation rule: If φ and ψ are formulas, then $\mathbf{X}\varphi$ and $\varphi\mathbf{U}\psi$ are formulas. Free and bound variables are defined in the obvious way. The *universal closure* of an LTL-FO formula $\varphi(\bar{x})$ with free variables \bar{x} is the formula $\forall \bar{x}\varphi(\bar{x})$. An LTL-FO sentence is the universal closure of an LTL-FO formula.

Let \mathcal{A} be an artifact schema. An LTL-FO sentence over \mathcal{A} is one where each FO component is over $\mathcal{DB}_{\mathcal{A}}$. The semantics of LTL-FO formulas is standard, and we describe it informally. Let $\Gamma = \langle \mathcal{A}, \Sigma \rangle$ be an artifact system, and $\forall \bar{x} \varphi(\bar{x})$ an LTL-FO sentence over \mathcal{A} . The artifact system Γ satisfies $\forall \bar{x} \varphi(\bar{x})$ iff every run of Γ satisfies it. Let $\rho = \{\rho_i\}_{i \geq 0}$ be a run of Γ , and let $\rho_{\geq j}$ denote $\{\rho_i\}_{i \geq j}$, for $j \geq 0$. Note that $\rho = \rho_{\geq 0}$. The run ρ satisfies $\forall \bar{x} \varphi(\bar{x})$ iff for each valuation ν of \bar{x} in D , $\rho_{\geq 0}$ satisfies $\varphi(\nu(\bar{x}))$. The latter is defined by structural induction on the formula. Satisfaction of an FO sentence ψ by ρ_i is defined in the obvious way. The semantics of Boolean operators is standard. The meaning of the temporal operators \mathbf{X} , \mathbf{U} is the following (where \models denotes satisfaction and $j \geq 0$):

- $\rho_{\geq j} \models \mathbf{X}\varphi$ iff $\rho_{\geq j+1} \models \varphi$,
- $\rho_{\geq j} \models \varphi \mathbf{U} \psi$ iff $\exists k \geq j$ such that $\rho_{\geq k} \models \psi$ and $\rho_{\geq l} \models \varphi$ for $j \leq l < k$.

Observe that the above temporal operators can simulate all commonly used operators, including \mathbf{F} (eventually), \mathbf{G} (always), and \mathbf{B} (before, which requires its first argument to hold before its second argument fails). Indeed, $\mathbf{F}\varphi \equiv \text{true} \mathbf{U} \varphi$, $\mathbf{G}\varphi \equiv \neg \mathbf{F}(\neg\varphi)$, and $\varphi \mathbf{B} \psi \equiv \neg(\neg\varphi \mathbf{U} \neg\psi)$. We use the above operators as shorthand in LTL-FO formulas whenever convenient.

Note that, as customary in verification, LTL-FO properties of artifact systems concern exclusively their infinite runs. Thus, blocking finite pre-runs are ignored. In particular, if an artifact system has only blocking pre-runs (so no proper run) then it vacuously satisfies all LTL-FO formulas. For this and other reasons, one may wish to know if, for a given artifact system (*i*) all of its pre-runs are blocking, or (*ii*) there exists a blocking pre-run. We consider decidability of these questions at the end of Section 3 (Corollary 3.4) and Section 4 (Corollary 4.3).

Example 2.11 We illustrate desirable properties for the artifact system Γ_{ex} in Example 2.9. These properties pertain to the global evolution of Γ_{ex} , as well as to the consistency of its specification. One such consistency property requires the state flags `in_process` and `done` in class `ORDER` to always be mutually exclusive:

$$\mathbf{G}(\neg(\text{in_process} \wedge \text{done})).$$

A more data-dependent consistency property requires line item quotes archived in state `li_quotes` of the `QUOTE` artifact to pertain only to tentative line items previously input by the customer (into attributes `li_prod` and `li_qty` of the `ORDER` artifact), and which underwent successful negotiation and approval. Successful negotiation occurs when `ask`, `bid` and `final_price` coincide and the `QUOTE` artifact is in state `archive`:

$$\begin{aligned} \forall pid, qty, price \mathbf{G} (& (\exists o, c, n R_O(o, c, n, pid, qty) \wedge \\ & \exists d, l, m R_Q(o, d, l, price, price, price, \text{"yes"}, pid, qty, m) \wedge \text{archive}) \\ & \mathbf{B} \\ & \neg \text{li_quotes}(pid, qty, price)). \end{aligned}$$

Notice the use of the *before* operator \mathbf{B} (requiring its first argument to hold before its second argument fails).

The following property is more semantic in nature, capturing part of the manufacturer's business model. It requires that if the customer's status is not *"preferred"* and the credit rating is worse than *"good"*, then before archiving a line item with final negotiated price lower than the manufacturer's desired price, explicit approval from a human executive must have been requested. We assume the following ordering on the constants indicating the credit rating: *"poor"* < *"fair"* < *"good"* < *"excellent"*.

$$\begin{aligned} \varphi_3 : \forall o, c, n, p, q, d, l, f, m, s, r \\ \mathbf{G} ((R_O(o, c, n, p, q) \wedge \text{in_process} \wedge R_Q(o, d, l, f, f, f, \omega, p, q, m) \wedge \text{negotiation} \wedge f < d \end{aligned}$$

$$\begin{aligned} & \wedge \text{CUSTOMER}(c, s, r) \wedge s \neq \text{"preferred"} \wedge r < \text{"good"} \\ \rightarrow & \text{approval_pending } \mathbf{B} \neg(\text{archive} \wedge \text{li_quotes}(p, q, f)) \\ &). \end{aligned}$$

Note that φ_3 involves both artifacts and the underlying database. If the negotiation process is described by service *refined_negotiation*, then the property happens to be satisfied: indeed, recall from its state rules that this service requests approval whenever the customer's status is not preferred and his credit rating is not excellent. In particular, this applies to customers whose rating is worse than good, according to the above ordering of credit ratings. \square

A detailed specification of all services involved in our running example can be found in Appendix A.

3 Decidable Verification

In this section we establish the main decidability result on verification of artifact systems.

It is easily seen that satisfaction of an LTL-FO formula by an artifact system is generally undecidable, using Trakhtenbrot's theorem. To obtain decidability, we introduce a restricted class of artifact systems and LTL-FO properties, called *guarded*. This is the analog to artifact systems of the input-boundedness restriction, first introduced by Spielmann in the context of ASM transducers [32], and subsequently used for Web service verification [10]. The guarded restriction requires a form of bounded quantification in formulas used in state update rules and LTL-FO properties. The pre-and-post conditions of services are restricted to be existential.

The guarded restriction is formulated as follows.

Definition 3.1 Let $\Gamma = \langle \mathcal{A}, \Sigma \rangle$ be an artifact system. The set of guarded FO formulas over \mathcal{A} is obtained by replacing in the definition of FO the quantification formation rule by the following:

- if φ is a formula, α is an atom using an attribute relation of some artifact of \mathcal{A} , $\bar{x} \subseteq \text{free}(\alpha)$, and $\bar{x} \cap \text{free}(\beta) = \emptyset$ for every state atom β in φ , then $\exists \bar{x}(\alpha \wedge \varphi)$ and $\forall \bar{x}(\alpha \rightarrow \varphi)$ are formulas.

An artifact system is guarded iff all formulas used in the state rules of its services are guarded, and all pre-and-post conditions are \exists^* FO formulas² in which all state atoms are ground. An LTL-FO sentence over \mathcal{A} is guarded iff all of its FO components are guarded.

Example 3.2 The artifact system Γ_{ex} in our running example is guarded. This includes the complete specification in Appendix A, which shows that the guardedness restriction still offers significant expressive power. For instance, notice that post-condition ψ^{rn} is an \exists^* FO formula with no non-ground state atoms (trivially so, as it mentions no state at all). In addition, in all state rules the quantified variables appear guarded by atoms using attribute relations R_O or R_Q . No quantified variables appear in any state atom because no such atoms are mentioned. See the state rules of service *include_line_item* in Appendix A for a less trivial example of guarded state rules. There, state atoms do occur in the rule body, but only with non-quantified variables. All properties listed in Example 2.11 are guarded.

For an example of an unguarded state rule, consider a relaxation of the insertion rule in S^{rn} demanding that executive approval be short-circuited and the `archive` flag be set for all preferred customers with better than fair rating:

$$\begin{aligned} \text{archive} \leftarrow & \exists o, d, l, a, f, app, p, q, m, r R_Q(o, d, l, a, a, f, app, q, m) \\ & \wedge \exists c, n R_O(o, c, n, p, q) \wedge \text{CUSTOMER}(c, \text{"preferred"}, r) \wedge r > \text{"fair"}. \end{aligned}$$

²Note that these formulas do not have to obey the restricted quantification formation rule.

The problem here is that quantified variable r does not appear in any attribute atom (indeed it cannot, since neither R_O nor R_Q have any rating attribute). While our undecidability results in Section 4 imply that not every unguarded rule or property can be equivalently rewritten into a guarded one, it is often possible to do so by slightly modifying the specification, such that it preserves the intended business process semantics, at the cost of widening the attribute relation. This happens to apply here. One can simply extend the attribute schema of R_O to include the customer’s rating, in addition to the originally included customer id. The rating attribute would be set at the same time as the customer id attribute. The latter is set by service *initiate_order* in Appendix A, using a guarded post-condition that would remain guarded after the proposed extension. \square

The main result on decidability of verification for artifact systems is the following.

Theorem 3.3 *It is decidable, given a guarded artifact system Γ and a guarded LTL-FO formula φ , whether every run of Γ satisfies φ . Furthermore, the complexity of the decision problem is PSPACE-complete for fixed arity schemas, and EXPSPACE otherwise.*

The main challenge in establishing the above result is that artifact systems are *infinite-state systems*, due to the presence of unbounded data. To deal with this, the key idea is to develop a concise, symbolic representation of equivalence classes of runs of Γ , called *pseudoruns*, that retain just the information needed to check satisfaction of φ , and can be generated in PSPACE without explicitly constructing any actual run or database. The high-level structure of the proof is similar to the one for decidability of verification for extended ASM transducers [10]. However, the result for the artifact model substantively extends previous ones in two significant ways: (i) runs of artifact systems may use infinitely many domain values (unlike extended ASM transducers where the domain of each run is restricted to the active domain of the finite database), and (ii) the underlying domain is ordered. These extensions require much more care in developing the pseudorun technique, and render the proof of decidability considerably more difficult. This proof is provided in Appendix B.

Finally, we consider the issue of blocking pre-runs. As remarked in Section 2, LTL-FO properties of artifact systems concern only their (infinite) runs and ignore blocking pre-runs. In particular, if an artifact system has only blocking pre-runs (so no proper run) then it vacuously satisfies all LTL-FO formulas. It therefore becomes of interest to know whether all pre-runs of an artifact system are blocking. Moreover, blocking may also be of interest for reasons specific to the application (see also discussion in Section 5). We can show the following.

Corollary 3.4 *It is decidable, given a guarded artifact system Γ , whether all pre-runs of Γ are blocking. Furthermore, the complexity is PSPACE for fixed-arity schemas, and EXPSPACE otherwise.*

Proof: The result follows immediately from Theorem 3.3. Indeed, all pre-runs of Γ are blocking iff Γ has no (infinite) runs iff $\Gamma \models \text{false}$. The latter is decidable with the stated complexities by Theorem 3.3. \square

One may also wish to know if a given artifact system has *some* blocking pre-run. Interestingly, this turns out to be undecidable for guarded artifact systems (see Corollary 4.3).

4 Boundaries of Decidability

In this section we consider several variations of our artifact model and relaxations of the guarded conditions and show that they lead to undecidability of verification. This suggests that the restrictions we presented in order to ensure decidability are quite tight. Due to space constraints, the presentation of the alternative models is informal.

Attributes versus states We first revisit the distinction between the attribute relation R and the state relation S in artifact classes $C = \langle R, S \rangle$. One might legitimately wonder if the separate treatment is relevant to verification. We next show that this is indeed the case. More precisely, consider a modification of the artifact model where the state S is treated in the same way as R , except that R holds a single tuple while S holds an entire relation. In particular, in the definition of a service using artifact class $C = \langle R, S \rangle$:

- the pre-and-post conditions of the service are \exists^* FO formulas using R , S and the database (with S -atoms no longer restricted to be ground as previously);
- as before, the initial value of S is empty;
- there are separate post-condition formulas ψ_R and ψ_S for R and S , defining their contents in the output (R consists, as before, of *one* arbitrary tuple satisfying ψ_R , while S consists of the *set* of tuples satisfying ψ_S , with active domain semantics to guarantee finiteness).

We refer to R as the *tuple attribute set* of C and to S as the *relational attribute set* of C . We refer to such artifact systems as *hybrid-attribute*. Note that, in this model, there are no longer separate state relations. Since there are no states, the guarded restriction on hybrid-attribute services now simply amounts to the \exists^* FO form of the pre-and-post conditions. The guarded restriction for LTL-FO properties remains unchanged. We can show the following (the proof is by reduction from the Post Correspondence Problem, see Appendix C).

Theorem 4.1 *It is undecidable, given a guarded hybrid-attribute artifact system Γ and guarded LTL-FO formula φ , whether $\Gamma \models \varphi$. Moreover, this holds even for singleton artifact systems whose relational attribute set consist of a single attribute, and for a fixed LTL-FO formula φ with no variables.*

Relaxing the guarded restrictions We now consider several relaxations of the guarded restrictions. It turns out that even very small such relaxations lead to undecidability of verification. Specifically, we consider the following: (i) allowing non-ground state atoms in pre-and-post conditions, (ii) allowing state projections in state update rules (a simple form of un-guarded quantification), (iii) allowing un-guarded quantification in the LTL-FO property, and (iv) extending LTL-FO with path quantifiers.

We can show that each of the relaxations (i)-(iv) leads to undecidability of verification. The proof of (i) is similar to that of Theorem 4.1. The proofs of (ii) and (iii) are by reduction from the implication problem for functional and inclusion dependencies, known to be undecidable [9]. The proof of (iv) is by reduction from validity of $\exists^*\forall^*$ FO sentences, also known to be undecidable [8]. The proofs of (ii)-(iv) can be easily adapted from analogous results obtained for extended ASM transducers [10]. We therefore omit the details.

Functional dependencies It is natural to ask whether the decidability of verification holds under the assumption that the database satisfies certain integrity constraints. Unfortunately, we show that even simple key dependencies lead to undecidability.

Theorem 4.2 *It is undecidable, given a guarded singleton artifact system Γ , a set of functional dependencies F over \mathcal{DB} , and a guarded LTL-FO sentence φ , whether $\rho \models \varphi$ for every run ρ of Γ on a database satisfying F . Moreover, this holds even if \mathcal{DB} consists of one binary and one unary relation, and F consists of a single key constraint on the binary relation.*

The proof is done by reduction from the PCP, similarly to Theorem 4.1 (details are omitted).

Existence of a blocking pre-run Recall the question raised in Section 2: does an artifact system have (i) *only* blocking pre-runs, or (ii) *some* blocking pre-run? We showed in Section 3 that (i) is decidable for guarded artifact systems (Corollary 3.4). Interestingly, (ii) turns out to be undecidable.

Corollary 4.3 *It is undecidable, given a guarded artifact system Γ , whether Γ has some blocking pre-run.*

The result is shown similarly to Theorems 4.1 and 4.2, by reduction from the PCP. The key idea is to first search for a match to the PCP (without assurance that the key dependency assumed in Theorem 4.2 is satisfied), and in case of success make continuance of the run contingent upon violation of the dependency. This reduces the existence of a solution to the PCP to the existence of a blocking pre-run.

Order versus successor Recall that decidability of verification holds under the assumption that the domain D is countable and equipped with a dense, total order \leq with no endpoints. If \leq is replaced by a successor relation on D , verification becomes undecidable. The proof is, again, by reduction from the PCP.

We note that it remains open whether verification remains decidable if some of the assumptions on \leq do not hold, for instance if \leq is not dense.

5 Further Applications

We next discuss several problems previously raised in the context of artifact systems, to which our results on verification can be beneficially applied.

Business rules We consider an extension of the artifact formalism in support of service reuse and customization. In practice, services are often provided by autonomous third-parties, who typically strive for wide applicability and impose as unrestrictive pre-conditions as possible. In contrast, the designer who incorporates third-party services into the business process often requires more control over when these services apply, in the form of more restrictive pre-conditions. Such additional control may also be needed to ensure compliance with business regulations formulated by third parties, independently of the specific application. To address such needs, [6] introduces *business rules*, which are conditions that can be super-imposed on the pre-conditions of existing services without changing their implementation.

We adopt the notion here and formalize it as follows. Given an artifact system $\Gamma = \langle \mathcal{A}, \Sigma \rangle$, we associate a set $\mathcal{B} = \{\beta_\sigma \mid \sigma \in \Sigma\}$ of business rules to the services in Σ . A *business rule* is a sentence in $\mathcal{L}_{\mathcal{A}}$, just like a service pre-condition.

For instance, we revisit our running example and assume that order shipment is modeled by the *ship* service, whose pre-condition only checks that the ORDER artifact is in state **done**. We also assume the existence of a *collect_payment* service, which applies when the ORDER is in state **done**. Finally, we assume that the ORDER artifact is extended with a **paid** boolean state flag which is set by the *collect_payment* service. Now we wish to super-impose the following business rule, which implements the policy that only platinum customers with excellent credit may get their order shipped before payment is received:

$$\beta_{ship} : \exists o, c, n, p, s, r \quad R_O(o, c, n, p, q) \wedge \text{CUSTOMER}(c, s, r) \wedge (s = \text{"platinum"} \wedge r = \text{"excellent"} \vee \text{paid}).$$

Verification under business rules The verification problem for artifact system Γ and property φ *under business rules* \mathcal{B} , denoted $\Gamma \models_{\mathcal{B}} \varphi$, means checking that every run of Γ' satisfies φ , where Γ' is obtained by adding each business rule as a pre-condition conjunct to its corresponding service in Γ . We say that a business rule is *guarded* if it is guarded when viewed as a service pre-condition. It follows immediately as a corollary of Theorem 3.3 that verification under \mathcal{B} is decidable if Γ , φ and all business rules in \mathcal{B} are guarded.

A related problem concerns *incremental* verification under business rules. Note that, if $\Gamma \models \varphi$, then $\Gamma \models_{\mathcal{B}} \varphi$. However, $\Gamma \not\models \varphi$ does *not* imply that $\Gamma \not\models_{\mathcal{B}} \varphi$. Thus, properties such as reachability of a configuration satisfying some desired property are not inherited when business rules are added. It is of interest whether such properties can be verified incrementally; however, we do not address this here.

Redundant business rules Towards streamlining the specification, a desirable goal is the removal of redundant business rules. This involves checking whether, given an artifact system $\Gamma = \langle \mathcal{A}, \Sigma \rangle$, a new business rule β associated to some service $\sigma \in \Sigma$ has any effect on Γ , i.e. excludes at least one of its runs. The latter problem amounts to verifying that at any point in a run of Γ , the pre-condition π of σ implies β : $\Gamma \models \mathbf{G}(\pi \rightarrow \beta)$. If Γ is guarded, and β is guarded in the sense of guarded FO components of LTL-FO properties, then, again as a corollary of Theorem 3.3, checking if β has an effect on Γ is decidable. Indeed, if π is guarded, then we have $\pi \doteq \exists \bar{x} f(\bar{x})$ with f a quantifier-free formula in $\mathcal{L}_{\mathcal{A}}$. Then

$$\Gamma \models \mathbf{G}(\exists \bar{x} f(\bar{x}) \rightarrow \beta) \text{ iff } \Gamma \models \mathbf{G}(\forall \bar{x} \neg f(\bar{x}) \vee \beta) \text{ iff } \Gamma \models \underbrace{\forall \bar{x} \mathbf{G}(\neg f(\bar{x}) \vee \beta)}_{\varphi}$$

where φ is a guarded LTL-FO property if β is. For example, β_{ship} above is guarded.

Redundant attributes Another design simplification consists of redundant attribute removal, a problem raised in [6]. We formulate this as follows. We would like to test whether there is a way to satisfy a property φ of runs without using one of the attributes, say a , of artifact A . Checking redundancy of a reduces to the following verification problem:

$$\Gamma \not\models \varphi \rightarrow \underbrace{\mathbf{F}(\exists \bar{x} \exists a R_A(\bar{x}, a) \wedge a \neq \omega)}_{\varphi'}$$

where we assume wlog that a is last in A 's attribute relation R_A . Recall from Section 2 the convention of representing undefined attributes using a constant ω . The argument of the temporal operator \mathbf{F} (*eventually*) checks that attribute a is defined. If φ is guarded and has no global variables (i.e. its FO components are all sentences), then φ' is a guarded LTL-FO property. Therefore Theorem 3.3 applies, yielding decidability.

Verifying termination properties Recall that our semantics of artifact systems and LTL-FO properties ignores blocking runs. However, in some applications, one would like to verify properties relating to termination. As discussed in Section 3, it is decidable if *all* pre-runs of an artifact system are blocking (Corollary 3.4). However, it may be desirable to verify more expressive properties involving blocking configurations. To this end, one can modify the semantics to render all runs infinite by repeating forever blocking configurations, whenever reached. It can be shown that our results continue to hold with this semantics. Note that one can state, within a guarded LTL-FO property, that a configuration of a guarded artifact system is blocking (all variables in negations of the \exists^* FO pre-conditions become globally quantified universally).

6 Conclusions

In this paper, we introduce the artifact system model, which formalizes a business process modeling paradigm that has recently attracted the attention of both the industrial and research communities. We study the problem of automatic verification of artifact systems, with the goal of increasing confidence in the correctness of such business processes.

All prior versions of the artifact model are inherently data-aware, being essentially evolved dataflow models. The version we consider extends prior models, taking significant additional steps towards data-awareness. It includes an underlying database which can be consulted by the services, and equips artifacts with updatable state relations. The service and property specifications allow sophisticated manipulation of data values via first-order formulae over the attributes and state of artifacts, the underlying database, and an infinite, ordered underlying domain. Data awareness raises a significant challenge compared to classical finite-state model checking, by turning artifact systems into infinite-state systems, whose verification problem is notoriously difficult.

We trace the boundaries of decidability for verification and we identify the guarded restriction, defining a practically appealing and fairly tight class of artifact systems and properties for which verification is decidable in PSPACE. This complexity is the best one can hope for, given that finite-state model checking is already PSPACE-complete. Our decidability result is significantly more difficult than the previous results of [32, 10] for ASM transducers and Web services, because each run is allowed to use infinitely many values from an underlying ordered domain. This extension is critical to the artifact framework, in order to adequately model arbitrary external input and partially specified processes given by pre- and post-conditions. Finally, we show that the verification techniques can also be leveraged to solve other static analysis tasks previously formulated for the artifact framework.

Our results are enabled by a mix of techniques from logic and model checking. We believe them to be of interest to the database, computer-aided verification, and business process communities.

References

- [1] S. Abiteboul, L. Herr, and J. V. den Bussche. Temporal versus first-order logic to query temporal databases. In *Proc. ACM PODS*, pages 49–57, 1996.
- [2] S. Abiteboul, L. Segoufin, and V. Vianu. Static analysis of Active XML systems. In *Proc. Intl. Symp. on Principles of Database Systems (PODS)*, pages 221–230, 2008.
- [3] S. Abiteboul, V. Vianu, B. Fordham, and Y. Yesha. Relational transducers for electronic commerce. *JCSS*, 61(2):236–269, 2000. Extended abstract in PODS 98.
- [4] K. Bhattacharya, N. S. Caswell, S. Kumaran, A. Nigam, and F. Y. Wu. Artifact-centered operational modeling: Lessons from customer engagements. *IBM Systems Journal*, 46(4):703–721, 2007.
- [5] K. Bhattacharya et al. A model-driven approach to industrializing discovery processes in pharmaceutical research. *IBM Systems Journal*, 44(1):145–162, 2005.
- [6] K. Bhattacharya, C. E. Gerede, R. Hull, R. Liu, and J. Su. Towards formal analysis of artifact-centric business process models. In *Proc. Int. Conf. on Business Process Management (BPM)*, pages 288–304, 2007.
- [7] K. Bhattacharya, R. Hull, and J. Su. A Data-centric Design Methodology for Business Processes. In J. Cardoso and W. van der Aalst, editors, *Handbook of Research on Business Process Management*. 2009. to appear.
- [8] E. Borger, E. Gradel, and Y. Gurevich. *The Classical Decision Problem*. Springer, 1997.
- [9] A. K. Chandra and M. Vardi. The implication problem for functional and inclusion dependencies is undecidable. *SIAM J. Comp.*, 14(3):671–677, 1985.

- [10] A. Deutsch, L. Sui, and V. Vianu. Specification and verification of data-driven web applications. *JCSS*, 73(3):442–474, 2007.
- [11] A. Deutsch, L. Sui, V. Vianu, and D. Zhou. Verification of communicating data-driven web services. In *Proc. Intl. Symp. on Principles of Database Systems (PODS)*, pages 90–99, 2006.
- [12] G. Dong, R. Hull, B. Kumar, J. Su, and G. Zhou. A framework for optimizing distributed workflow executions. In *Proc. Intl. Workshop on Database Programming Languages (DBPL)*, pages 152–167, 1999.
- [13] E. A. Emerson. Temporal and modal logic. In J. V. Leeuwen, editor, *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*, pages 995–1072. North-Holland Pub. Co./MIT Press, 1990.
- [14] M. R. Garey and D. S. Johnson. *Computers and Intractability*. Freeman, 1979.
- [15] C. E. Gerede, K. Bhattacharya, and J. Su. Static analysis of business artifact-centric operational models. In *IEEE International Conference on Service-Oriented Computing and Applications*, 2007.
- [16] C. E. Gerede and J. Su. Specification and verification of artifact behaviors in business process models. In *Proceedings of 5th International Conference on Service-Oriented Computing (ICSOC)*, Vienna, Austria, September 2007.
- [17] R. Glushko and T. McGrath. *Document Engineering: Analyzing and Designing Documents for Business Informatics and Web Services*. MIT Press, Cambridge, MA, 2005.
- [18] R. Hull, F. Llirbat, B. Kumar, G. Zhou, G. Dong, and J. Su. Optimization techniques for data-intensive decision flows. In *Proc. IEEE Intl. Conf. on Data Engineering (ICDE)*, pages 281–292, 2000.
- [19] R. Hull, F. Llirbat, E. Simon, J. Su, G. Dong, B. Kumar, and G. Zhou. Declarative workflows that support easy modification and dynamic browsing. In *Proc. Int. Joint Conf. on Work Activities Coordination and Collaboration*, 1999.
- [20] S. Kumaran, R. Liu, and F. Y. Wu. On the duality of information-centric and activity-centric models of business processes. In *Proc. Intl. Conf. on Advanced Information Systems Engineering (CAISE)*, 2008.
- [21] S. Kumaran, P. Nandi, T. Heath, K. Bhaskaran, and R. Das. ADoc-oriented programming. In *Symp. on Applications and the Internet (SAINT)*, pages 334–343, 2003.
- [22] J. Küster, K. Ryndina, and H. Gall. Generation of BPM for object life cycle compliance. In *Proceedings of 5th International Conference on Business Process Management (BPM)*, 2007.
- [23] R. Liu, K. Bhattacharya, and F. Y. Wu. Modeling business contexture and behavior using business artifacts. In *CAiSE*, volume 4495 of *LNCS*, 2007.
- [24] D. Martin et al. OWL-S: Semantic markup for web services, W3C Member Submission, November 2003.
- [25] S. A. McIlraith, T. C. Son, and H. Zeng. Semantic web services. *IEEE Intelligent Systems*, 16(2):46–53, 2001.
- [26] P. Nandi and S. Kumaran. Adaptive business objects – a new component model for business integration. In *Proc. Intl. Conf. on Enterprise Information Systems*, pages 179–188, 2005.
- [27] S. Narayanan and S. McIlraith. Simulation, verification and automated composition of web services. In *Intl. World Wide Web Conf. (WWW2002)*, 2002.

- [28] A. Nigam and N. S. Caswell. Business artifacts: An approach to operational specification. *IBM Systems Journal*, 42(3):428–445, 2003.
- [29] E. L. Post. Recursive unsolvability of a problem of Thue. *J. of Symbolic Logic*, 12:1–11, 1947.
- [30] A. P. Sistla, M. Y. Vardi, and P. Wolper. The complementation problem for Büchi automata with applications to temporal logic. *Theoretical Computer Science*, 49:217–237, 1987.
- [31] M. Spielmann. Abstract State Machines: Verification problems and complexity. Ph.D. thesis, RWTH Aachen, 2000.
- [32] M. Spielmann. Verification of relational transducers for electronic commerce. *JCSS.*, 66(1):40–65, 2003. Extended abstract in PODS 2000.
- [33] M. Y. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *Symp. on Logic in Computer Science*, 1986.
- [34] J. Wang and A. Kumar. A framework for document-driven workflow systems. In *Business Process Management*, pages 285–301, 2005.

Appendix

A Running Example

We present here our full running example. We model a scenario where a manufacturer fills customer purchase orders, negotiating the price of each line item on a case-by-case basis. We focus on two artifacts manipulated by the negotiation process, ORDER and QUOTE. During the workflow, the customer repeatedly adds new line items into (or updates existing ones in) the purchase order modeled by the ORDER artifact. Each line item specifies a product and its quantity. Line items are first tentatively filled in the ORDER attributes li_prod and li_qty . Every tentative line item spawns a negotiation process, in which manufacturer and customer complete rounds of declaring ask and bid prices, until agreement is reached or the negotiation fails. The prices at every round are stored in the QUOTE artifact, which also holds the manufacturer’s initially desired price, the lowest bid he is willing to entertain, and the final negotiated price. Once the negotiation on a tentative line item succeeds, its outcome is scrutinized by a human executive working for the manufacturer. Upon the executive’s approval, the line item is included into the purchase order (by insertion into the ORDER state $line_items$), and the final price is archived (in the QUOTE state li_quotes). During the negotiation, the manufacturer consults an underlying database, which lists information about available products (e.g. manufacturing cost) and about customers (e.g. credit rating and status).

The corresponding artifact system $\Gamma_{ex} = \langle \mathcal{A}, \Sigma \rangle$ is formally described below. As a font convention, we use R to refer to an artifact’s attribute relation and S for state relations.

The artifact schema is $\mathcal{A} = \langle ORDER, QUOTE, \mathcal{DB} \rangle$, detailed as follows.

1. $\mathcal{DB} = \langle PRODUCT, CUSTOMER \rangle$ is the database schema, where:

- $PRODUCT(prod_id, manufacturing_cost, min_order_qty)$ is the relation containing products and production information, and
- $CUSTOMER(customer_id, status, credit_rating)$ contains information about customers, such as customer status and credit rating.

2. $ORDER = \langle R_O, \underbrace{line_items, in_process, done}_{S_O} \rangle$

is the artifact class containing the information about a customer’s order.

- $R_O(order\#, customer_id, need_by, li_prod, li_qty)$ is the attribute relation holding the order number, the identifier of the customer who placed the order, the day it is needed by. The role of attributes li_prod and li_qty is described below.
- $line_items(prod_id, qty)$ is a state relation that acts as a “shopping cart” holding the collection of line items requested so far.
- $in_process$ and $done$ are nullary state relations (boolean flags) keeping track of the stage the artifact is in.³

The intention is that, in stage $in_process$, the customer repeatedly updates the shopping cart by specifying an individual, tentative line item described by attributes li_prod and li_qty . Subsequently, this line item is

³Artifact class ORDER illustrates an extension of Definition 2.1 that allows several state relations. This extension is for convenience only: it is easy to show that it provides no additional expressive power and preserves our decidability results. Indeed, given artifact system Γ with multiple states per artifact and LTL-FO sentence φ , we can construct in polynomial time artifact system Γ' and sentence φ' such that $\Gamma \models \varphi$ iff $\Gamma' \models \varphi'$. Moreover, if Γ and φ are guarded, then so are Γ' and φ' . The proof is similar to the proof of Lemma B.1.

inserted into, deleted from, or replaces in `line_items` an item with the same `prod_id`, provided the price negotiation succeeds. When the customer completes the purchase order, the ORDER artifact transitions to stage `done`.

3. $QUOTE = \langle R_Q, \underbrace{\text{li_quotes, idle, desired_price_calc, negotiation, approval_pending, archive}}_{S_Q} \rangle$

is the artifact class modeling quotes, with:

- $R_Q(\text{order}\#, \text{desired_price}, \text{lowest_acceptable_price}, \text{ask}, \text{bid}, \text{final_price}, \text{approved}, \text{li_prod}, \text{li_qty}, \text{manufacturing_cost})$
is the attribute relation.
- $\text{li_quotes}(\text{prod_id}, \text{qty}, \text{price})$
is a state relation holding the final negotiated price quotes for the line items in the corresponding ORDER artifact.
- `idle`, `desired_price_calc`, `negotiation`, `approval_pending`, `archive`
are nullary state relations.

When inactive, the QUOTE artifact is in state `idle`, but moves to `desired_price_calc` as soon as the customer fills in the product id and quantity of a line item. In this stage, `desired_price` attribute is set (from the manufacturer's point of view), possibly taking into account the `need_by` date attribute in the corresponding ORDER artifact and the manufacturing cost listed in the PRODUCT database. During the ensuing `negotiation` stage, the ask and bid prices are repeatedly set (in attribute `ask` by the manufacturer, respectively `bid` by the customer) until a final price is established and recorded in attribute `final_price`, or the negotiation fails. Final prices may require approval by a human executive to whom the negotiator reports. While approval is awaited, the QUOTE artifact is in stage `approval_pending`. Approval is granted by setting boolean attribute `approved`. Approved final prices are then archived in state relation `li_quotes` (while the QUOTE artifact is in stage `archive`).

The operations allowed on artifacts are modeled by the set Σ of available services, summarized below.

- Service $\text{initiate_order} = \langle \pi^{io}, \psi^{io}, \mathcal{S}^{io} \rangle$ initializes the ORDER artifact, modeling the input of the order number and customer id by the manufacturer, and the need-by date by the customer.
- Service $\text{add_or_modify_line_item} = \langle \pi^{am}, \psi^{am}, \mathcal{S}^{am} \rangle$ models the customer's choice of a tentative line item to be added to the purchase order, or to replace another line item for the same product. The service records this choice in attributes `li_prod` and `li_qty` of the ORDER artifact, and initializes the QUOTE artifact in view of the upcoming negotiation. This involves copying the order number and line item information from ORDER to QUOTE, and filling the QUOTE's `manufacturing_cost` attribute with the corresponding value looked up in the PRODUCT database.
- Service $\text{include_line_item} = \langle \pi^{il}, \psi^{il}, \mathcal{S}^{il} \rangle$ includes into the purchase order the current tentative line item, by storing it in ORDER state `line_items`. The corresponding final negotiated price is archived in QUOTE state `li_quotes`.
- Service $\text{commit_order} = \langle \pi^{co}, \psi^{co}, \mathcal{S}^{co} \rangle$ simply switches the ORDER artifact to the `done` stage, which disables any further line item modifications. This service models the customer's non-deterministic decision to finalize the purchase order.

- Service $set_quote_interval = \langle \pi^{sq}, \psi^{sq}, \mathcal{S}^{sq} \rangle$ sets the *desired_price* and *lowest_acceptable_price* attributes of the QUOTE artifact to frame the subsequent negotiation. This service abstracts a complex sub-task, possibly taking into account the ORDER's *need_by* attribute, the manufacturer's desired profit margin, the customer's status, and input from a human manager.
- Service $quote_approval = \langle \pi^{qa}, \psi^{qa}, \mathcal{S}^{qa} \rangle$ models the human supervisor who reviews the quote on behalf of the manufacturer. The process is a black box, about which is only known that it switches the QUOTE artifact to *archive* stage, and it sets the *approved* attribute to either "yes" or "no".

To showcase the artifact model's natural ability to specify processes even partially, we describe the negotiation process at two levels of abstraction.

- In a first, coarser cut, the process is abstracted as service $abstract_negotiation = \langle \pi^{an}, \psi^{an}, \mathcal{S}^{an} \rangle$ about which we only know that the final price is reached when the ask and bid prices coincide, and that it is guaranteed to lie between the allowed margins stored in attributes *desired_price* and *lowest_acceptable_price* of artifact QUOTE.
- Alternatively, we use service $refined_negotiation = \langle \pi^{rn}, \psi^{rn}, \mathcal{S}^{rn} \rangle$ to refine the negotiation process all the way to the level of individual negotiation rounds, each of which sets the current ask and bid prices.

Conventions We adopt the following conventions:

- (i) We model uninitialized attributes by setting them to the reserved constant ω .
- (ii) We model Boolean states by nullary state relations, and drop the parentheses from atoms using them: $\mathbf{S}()$ becomes \mathbf{S} . We assume the usual encoding of *true* as the singleton nullary relation, and *false* as the empty nullary relation. In particular, all Boolean states are initially *false* (since all state relations are initially empty).
- (iii) For convenience, we use the following syntactic sugar for post-conditions: we write post-conditions as non-Horn rules $h(\bar{x}) := b(\bar{y})$ where the *head* h is a conjunction of atoms over attribute relations in \mathcal{A} , with variables \bar{x} , and the *body* b is a formula in $\mathcal{L}_{\mathcal{A}}$ with free variables \bar{y} , where $\bar{y} \subseteq \bar{x}$. The semantics is that whenever $A \xrightarrow{\sigma} A'$ holds, $A' \models h(\bar{x} \leftarrow \bar{u})$ for some tuple \bar{u} , and $A \models b(\bar{y} \leftarrow \bar{u}|\bar{y})$. Moreover, artifact relations not mentioned in h remain unchanged. Clearly, this syntactic sugar can be simulated by the official post-conditions, and conversely.

Service $initiate_order = \langle \pi^{io}, \psi^{io}, \mathcal{S}^{io} \rangle$ initializes the ORDER artifact, where:

- $\pi^{io} \doteq \neg in_process$,
i.e. the service applies when the ORDER artifact is not used to process another order;
- The post-condition ψ^{io} given by

$$R_O(o, c, n, \omega, \omega) := o \neq \omega \wedge n \neq \omega \wedge \exists r, s \text{ CUSTOMER}(c, s, r)$$

guarantees that the attributes *order#*, *customer_id* and *need_by* are initialized (set distinct from ω). By the semantics of post-conditions, the pick of o, c, n is non-deterministic. The pick of n models the customer's input, while that of o models the assignment of an order number by the manufacturer. Note that no further constraints are imposed on these values, they are simply picked from the infinite domain. In contrast, the customer c must be one of the existing customers listed in the database relation CUSTOMER. The pick of c also models the manufacturer's input.

The tentative line item's price p and quantity q are left uninitialized (they equal ω) and will be set by the customer during an activity modeled by service $add_or_modify_line_item$ below.

- The state rules in \mathcal{S}^{io} include the following:
 - `in_process` $\leftarrow true$,
an insertion rule that sets the `in_process` boolean flag.
 - `¬done` $\leftarrow true$,
a deletion rule that resets boolean state flag `done`, ensuring it is mutually exclusive with `in_process`.
(it is the responsibility of all other services operating on `ORDER` to keep them so
– see *add_or_modify_line_item* below).

No rule refers to state relation `line_items`, as no line item exists yet.

Service *add_or_modify_line_item* models the customer’s choice of a tentative line item to be added to the purchase order, or to replace another line item for the same product. The service affects both the `ORDER` and the `QUOTE` artifact, and it is given as $\langle \pi^{am}, \psi^{am}, \mathcal{S}^{am} \rangle$, where:

- $\pi^{am} \doteq \exists o, c, n \text{ in_process} \wedge R_O(o, c, n, \omega, \omega) \wedge \text{idle} \wedge R_Q(\omega, \omega, \omega, \omega, \omega, \omega, \omega, \omega, \omega, \omega)$,
i.e. the service applies only if the `ORDER` artifact is in the `in_process` stage, no other line item is currently being processed, and if the `QUOTE` artifact is currently unused (in the `idle` stage, with all attributes uninitialized).
- The post-condition ψ^{am} is

$$R_O(o, c, n, p, q) \wedge R_Q(o, \omega, \omega, \omega, \omega, \omega, p, q, m) := \\ \exists q' R_O(o, c, n, \omega, \omega) \wedge p \neq \omega \wedge q \neq \omega \wedge q \geq q' \wedge \text{PRODUCT}(p, m, q')$$

Note that the customer’s input of product id p and quantity q is modeled as a non-deterministic pick from the infinite domain. The picked p must appear in the `PRODUCT` catalog stored in the database. The quantity q is less restricted: we only know that it is defined ($q \neq \omega$) and, reflecting the manufacturer’s policy, it exceeds the minimum-order quantity q' listed in the `PRODUCT` catalog.

According to the post-condition, the service reacts as follows to the customer’s input of the tentative line item. It stores the values p and q into the attributes `li_prod`, `li_qty` of the `ORDER` artifact. Note that attributes `order#`, `customer_id` and `need_by` remain unchanged. The service also initializes the `order#` attribute of the `QUOTE` artifact to refer to the corresponding order, and also stores in it p , q and the manufacturing cost m for product p , which is looked up in the database catalog `PRODUCT`. The `QUOTE` artifact’s remaining attributes are left undefined, to be set during negotiation.

- \mathcal{S}^{am} contains no `ORDER` state rule as the order’s state is left unchanged. It contains the following state rules that move the `QUOTE` artifact to the `desired_price_calc` stage, which enables the sub-task of quote negotiation: insertion rule

$$\text{desired_price_calc} \leftarrow true$$

and deletion rule

$$\neg \text{idle} \leftarrow true.$$

Service *include_line_item* = $\langle \pi^{il}, \psi^{il}, \mathcal{S}^{il} \rangle$ includes into the purchase order the current tentative line item, by storing it in `ORDER` state `line_items`. The corresponding final negotiated price is archived in `QUOTE` state `li_quotes`. We have:

- The pre-condition

$$\pi^{il} \doteq \text{in_process} \wedge \exists o, c, n, p, q R_O(o, c, n, p, q) \wedge o \neq \omega \wedge c \neq \omega \wedge n \neq \omega \wedge p \neq \omega \wedge q \neq \omega \wedge \\ \exists d, l, f, m R_Q(o, d, l, f, f, f, \text{"yes"}, p, q, m) \wedge \text{archive}$$

ensures that the service applies only if a current line item p, q exists, the ORDER artifact is in stage `in_process`, and the QUOTE artifact lists a successful and approved negotiation for this line item and order (notice the common occurrence of o, p, q in both the QUOTE and the ORDER atoms). Successful negotiation occurs when the ask, bid and final price coincide, and the artifact is in state `archive`. The final price is approved when the *approved* attribute is set to “yes”.

- The post-condition ψ^{il} given by

$$R_O(o, c, n, \omega, \omega) \wedge R_Q(\omega, \omega, \omega, \omega, \omega, \omega, \omega, \omega, \omega, \omega) := \exists p', q' R_O(o, c, n, p', q')$$

guarantees that, regardless of the current value p', q' of the line item, in the successor ORDER artifact the values are reset to undefined (ω), to make room for the next tentative line item. Notice that the ORDER attributes *order#*, *customer_id* and *need_by* are preserved. The QUOTE attributes are all reset, in preparation for the next negotiation.

- The state rules in \mathcal{S}^{il} include the following:

- The state insertion rule

$$\text{line_items}(p, q) \leftarrow \exists o, c, n R_O(o, c, n, p, q)$$

operates on artifact ORDER, inserting the values of attributes *li_prod* and *li_qty* into state relation `line_items`.

The state deletion rule

$$\neg \text{line_items}(p, q) \leftarrow \exists o, c, n, q' R_O(o, c, n, p, q') \wedge \text{line_items}(p, q)$$

deletes any other entry pertaining to the same product p (if any). Recall that, according to the possible successor definition, if state `line_items` already contains an entry for product p , the combined effect of the insertion and deletion rule is that of *updating* the quantity of product p to the latest customer-provided (and successfully negotiated) value. If no prior entry for p exists, then the deletion rule has no effect.

- The final negotiated price for this line item is archived in QUOTE state `li_quotes` by the following insertion rule:

$$\text{li_quotes}(p, q, f) \leftarrow \exists o, d, l, m R_Q(o, d, l, f, f, f, \text{"yes"}, p, q, m).$$

The following insertion and deletion rules move the QUOTE artifact to state `idle`, signaling its availability for a new negotiation sub-task:

$$\text{idle} \leftarrow \text{true} \text{ and} \\ \neg \text{archive} \leftarrow \text{true}.$$

Service *commit_order* = $\langle \pi^{co}, \psi^{co}, \mathcal{S}^{co} \rangle$ simply switches the ORDER artifact to the `done` stage, which disables any further line item modifications. This service models the customer’s non-deterministic decision to finalize the purchase order.

- $\pi^{co} \doteq \text{in_process} \wedge \exists o, c, n, p, q R_O(o, c, n, p, q) \wedge p = \omega \wedge q = \omega$,
i.e. the full order can be committed only if no tentative line item is still being processed (which would make $p \neq \omega, q \neq \omega$).

- ψ^{co} is given by

$$R_O(o, c, n, p, q) := R_O(o, c, n, p, q)$$

i.e. the artifact's attributes do not change.

- \mathcal{S}^{co} contains only the rules
 $\text{in_process} \leftarrow \text{false}$ and
 $\text{done} \leftarrow \text{true}$.

The following services model the negotiation process.

Service *set_quote_interval* sets the *desired_price* and *lowest_acceptable_price* attributes of the QUOTE artifact to frame the subsequent negotiation. This service abstracts a complex sub-task, possibly taking into account the ORDER's *need_by* attribute, the manufacturer's desired profit margin, the customer's status, and input from a human manager.

set_quote_interval = $\langle \pi^{sq}, \psi^{sq}, \mathcal{S}^{sq} \rangle$, where:

- $\pi^{sq} \doteq \text{desired_price_calc}$.
- Post-condition ψ^{sq} given by

$$R_Q(o, d, l, d, \omega, \omega, \omega, p, q, m) := d \neq \omega \wedge l \neq \omega \wedge d \geq l \geq m \wedge R_Q(o, \omega, \omega, \omega, \omega, \omega, \omega, p, q, m)$$

models only what is known about the quote generation procedure viewed as a black box: namely that the desired price is higher than the lowest acceptable one, which in turn exceeds the manufacturing cost. It also sets the initial asking price to the desired price in preparation for the negotiation stage.

- The rules in \mathcal{S}^{sq} simply switch the artifact to the negotiation stage, and are omitted.

To showcase the artifact model's natural ability to specify processes even partially, we describe the negotiation process at two levels of abstraction.

In a first, coarser cut, the process is abstracted as service *abstract_negotiation* = $\langle \pi^{an}, \psi^{an}, \mathcal{S}^{an} \rangle$ about which we only know that the final price is reached when the ask and bid prices coincide, and that it is guaranteed to lie between the allowed margins stored in attributes *desired_price* and *lowest_acceptable_price* of artifact QUOTE.

- $\pi^{an} = \text{negotiation}$,
since the process can only start when the QUOTE artifact is ready, which is signaled by setting this state flag.
- Post-condition ψ^{an} , given as

$$R_Q(o, d, l, f, f, f, app, p, q, m) := \exists a', b', f' R_Q(o, d, l, a', b', f', app, p, q, m) \wedge l \leq f \leq d$$

guarantees that the final price f agrees with the final ask and bid prices regardless of their initial values a', b' , and that f lies between the desired price d and the lowest acceptable price l .

- We omit the rules in \mathcal{S}^{an} , which move the artifact to state *approval_pending*.

Alternatively, we can refine the negotiation process all the way to the level of individual negotiation rounds, each of which sets the current ask and bid prices. A post-condition ensures that the negotiation is well-formed, i.e. the bid never exceeds the ask price, and that across rounds, asking prices never increase, while bids never decrease. The negotiation is successful when ask and bid prices agree, at which time the QUOTE artifact moves to the `approval_pending` state, the `final_price` attribute is set, and no further rounds are conducted.

Service `refined_negotiation` = $\langle \pi^{rn}, \psi^{rn}, \mathcal{S}^{rn} \rangle$ is described as follows:

- π^{rn} = **negotiation**
ensures that the service applies only as long as the boolean state flag **negotiation** is set in the QUOTE artifact.
- Post-condition ψ^{rn} is given as

$$R_Q(o, d, l, a, b, f, app, p, q, m) := \\ (\exists a', b' R_Q(o, d, l, a', b', \omega, app, p, q, m) \wedge a' \neq b' \wedge l \leq a \leq a' \wedge b' \leq b \wedge f = \omega) \\ \vee \\ (R_Q(o, d, l, a, b, \omega, app, p, q, m) \wedge a = b = f).$$

According to the first disjunct, the negotiation is well-formed, i.e. the bid never exceeds the ask price, and in each round, asking prices a never increase while bids b never decrease. Moreover, as long as ask and bid price differ, the final price remains undefined (equal to ω). Notice that the values of a and b are otherwise unconstrained, being simply drawn from the infinite domain. This models their external input by the manufacturer, respectively customer. The second disjunct states that once ask price a and bid price b coincide, the final price f is automatically set to the common value.

- \mathcal{S}^{rn} contains rules that, upon detecting successful negotiation, switch the QUOTE artifact to stage **approval_pending** if the customer does not enjoy preferred status with excellent credit. If he does, then the approval is short-circuited and the QUOTE goes to stage **archive**. The negotiation is successful when ask and bid prices agree.

$$\begin{aligned} \text{approval_pending} &\leftarrow \exists o, d, l, a, f, app, p, q, m R_Q(o, d, l, a, a, f, app, q, m) \\ &\quad \wedge \exists c, n R_O(o, c, n, p, q) \wedge \neg \text{CUSTOMER}(c, \text{"preferred"}, \text{"excellent"}) \\ \text{archive} &\leftarrow \exists o, d, l, a, f, app, p, q, m R_Q(o, d, l, a, a, f, app, q, m) \\ &\quad \wedge \exists c, n R_O(o, c, n, p, q) \wedge \text{CUSTOMER}(c, \text{"preferred"}, \text{"excellent"}) \\ \neg \text{negotiation} &\leftarrow \exists o, d, l, a, f, app, p, q, m R_Q(o, d, l, a, a, f, app, p, q, m) \end{aligned}$$

Note that state flag **negotiation** must be set before the state update rules execute (since pre-condition π^{rn} is satisfied). If neither of the state rule bodies is satisfied, then according to the possible successor semantics, the **negotiation** flag remains set, enabling another negotiation round.

Finally, service `quote_approval` = $\langle \pi^{qa}, \psi^{qa}, \mathcal{S}^{qa} \rangle$ models the human supervisor who reviews the quote on behalf of the manufacturer. The process is a black box, about which is only known that it switches the QUOTE artifact to **archive** stage, and it sets the `approved` attribute to either `"yes"` or `"no"`.

- $\pi^{qa} \doteq$ **approval_pending**.
- ψ^{qa} is given by

$$R_Q(o, d, l, f, f, f, app, p, q, m) := R_Q(o, d, l, f, f, f, \omega, p, q, m) \wedge (app = \text{"yes"} \vee app = \text{"no"}).$$

- \mathcal{S}^{qa} comprises the state rules that switch the artifact to stage `archive`, and are omitted.

We illustrate desirable properties for Γ_{ex} , which pertain to its global evolution, as well as to the consistency of the specification.

One such consistency property requires the state flags `in_process` and `done` in class `ORDER` to always be mutually exclusive:

$$\mathbf{G}(\neg(\text{in_process} \wedge \text{done})).$$

A more data-centric consistency property requires line item quotes archived in state `li_quotes` of the `QUOTE` artifact to pertain only to tentative line items previously input by the customer (into attributes `li_prod` and `li_qty` of the `ORDER` artifact), and which underwent successful negotiation and approval. Successful negotiation occurs when `ask`, `bid` and final price coincide and the `QUOTE` artifact is in state `archive`:

$$\begin{aligned} \forall pid, qty, price \mathbf{G} (& (\exists o, c, n R_O(o, c, n, pid, qty) \wedge \\ & \exists d, l, m R_Q(o, d, l, price, price, price, "yes", pid, qty, m) \wedge \text{archive}) \\ & \mathbf{B} \\ & \neg \text{li_quotes}(pid, qty, price)). \end{aligned}$$

Notice the use of the *before* operator \mathbf{B} (requiring its first argument to hold before its second argument fails).

The following property is more semantic in nature, capturing part of the manufacturer's business model. It requires that if the customer's status is not "*preferred*" and the credit rating is worse than "*good*", then before archiving a line item with final negotiated price lower than the manufacturer's desired price, explicit approval from a human executive must have been requested. We assume the following ordering on the constants indicating the credit rating: "*poor*" < "*fair*" < "*good*" < "*excellent*".

$$\begin{aligned} \varphi_3 : \forall o, c, n, p, q, d, l, f, m, s, r \\ \mathbf{G} (& (R_O(o, c, n, p, q) \wedge \text{in_process} \wedge R_Q(o, d, l, f, f, \omega, p, q, m) \wedge \text{negotiation} \wedge f < d \\ & \wedge \text{CUSTOMER}(c, s, r) \wedge s \neq \text{"preferred"} \wedge r < \text{"good"}) \\ & \rightarrow \text{approval_pending} \mathbf{B} \neg(\text{archive} \wedge \text{li_quotes}(p, q, f)) \\ &). \end{aligned}$$

Note that φ_3 involves both artifacts and the underlying database. If the negotiation process is described by service *refined_negotiation*, then the property happens to be satisfied: indeed, recall from its state rules that this service requests approval whenever the customer's status is not preferred and his credit rating is not excellent. In particular, this applies to customers whose rating is worse than good, according to the above ordering of credit ratings.

B Proof of main decidability result

We outline in this section the main steps in the proof of Theorem 3.3.

To begin, we show that it is enough to focus on artifact systems consisting of a single service operating on a single artifact; we call such systems singleton artifact systems.

Lemma B.1 *For each artifact system Γ and LTL-FO sentence φ , there exist a singleton artifact system Γ_s and LTL-FO sentence φ_s , computable in polynomial time from Γ and φ , such that $\Gamma \models \varphi$ iff $\Gamma_s \models \varphi_s$. Furthermore, if Γ and φ are guarded, then so are Γ_s and φ_s .*

Proof: We sketch the construction of Γ_s and φ_s in two stages. First, we show how one can derive from $\Gamma = \langle \mathcal{A}, \Sigma \rangle$ and $\varphi \in \mathcal{L}_{\mathcal{A}}$ an artifact system $\Gamma_r = \langle \mathcal{A}_r, \Sigma_r \rangle$ containing only one artifact class \mathcal{C}_r^1 , and an LTL-FO

formula $\varphi_r \in \mathcal{L}_{\mathcal{A}_r}$ such that $\Gamma \models \varphi \Leftrightarrow \Gamma_r \models \varphi_r$. Then, we show how Γ_s and φ_s can be obtained by “merging” the services of Γ_r into a single service and modifying the property accordingly.

For the first step, assume that Γ contains artifact classes $\mathcal{C}_i = \langle R_i, S_i \rangle$ ($i = 1, \dots, n$) and define $\mathcal{C}_1^r = \langle R_1^r, S_1^r \rangle$ as the only artifact class in Γ_r , with $a(R_1^r) = \sum_{i=1}^n a(R_i)$ and $a(S_1^r) = \sum_{i=1}^n a(S_i)$. The idea is to use R_1^r (respectively S_1^r) to store the content of R_1, \dots, R_n (S_1, \dots, S_n). In particular, the columns of R_1^r (S_1^r) are partitioned into n intervals, and the i -th partition is used to store tuples from R_i^r (S_i^r). Each service of Γ has a counterpart in Γ_r that accesses R_1^r . Of course, we need to modify the services’ pre/post-conditions and state rules so that they apply to only one attribute and state relation. This can be achieved by means of a syntactic transformation ζ that replaces each atom referring to R_i (S_i) with an FO formula that “retrieves” R_i (S_i) from R_1^r (S_1^r), by projecting out all components irrelevant to R_i (S_i requires an encoding using constants to avoid quantification violating the guarded restriction, but the approach is basically the same). Similar manipulations are required for state rules.

For the second step, suppose we constructed Γ_r and φ_r as above, with a single artifact class but multiple services acting on that class. The main difficulty in constructing Γ_s (with a single service σ_s) is to simulate the nondeterminism on service selection. To this end, we exploit the nondeterminism of postconditions. The idea is to extend relation R_1^r with an additional attribute, say *Service*, whose value identifies the next service of Γ_r to be simulated by σ_s . The post-condition of σ_s allows as the next value of *Service* any identifier of a service of Γ_r . The run of Γ_s blocks if, at any step, the precondition of the service identified by *Service* does not hold (but this does not pose a problem, since blocking runs do not affect LTL properties). The formula φ_r is straightforwardly modified to φ_s so that $\Gamma_r \models \varphi_r$ iff $\Gamma_s \models \varphi_s$ (the details are omitted). The entire construction preserves the guarded restriction and is done in polynomial time. \square

Let $\Gamma = \langle \mathcal{A}, \Sigma \rangle$ be a singleton artifact system. For notational simplicity, we henceforth specify such a system as a tuple $\langle \mathcal{DB}, R, S, \pi, \psi, \mathcal{S} \rangle$ where \mathcal{DB} is the database schema of Γ , $\langle R, S \rangle$ is the only artifact class, and π , ψ , and \mathcal{S} are the pre-condition, post-condition, and state rules of the service. We represent an instance of Γ by a tuple $\langle \text{DB}, R, S \rangle$ where DB is an instance of \mathcal{DB} , R is a relation over R containing a single tuple, and S is a relation over S .

We assume familiarity with model checking of LTL formulas using Büchi automata [33]. We recall that a Büchi automaton is a non-deterministic finite-state machine accepting infinite words (ω -words) over a finite alphabet by going infinitely many times through an accepting state. Given a propositional LTL-FO formula φ over a set P of propositions, there is a Büchi automaton A_φ accepting precisely the ω -words over alphabet 2^P that satisfy φ . Moreover, A_φ has exponentially many states (in $|\varphi|$). Fortunately A_φ need not be constructed explicitly to perform model checking: indeed, there exists a non-deterministic PSPACE algorithm that, given as input a state of A_φ and a set in 2^P , produces as output a next possible state of A_φ [30]. We will make use of this in our algorithm.

Consider a guarded singleton artifact system Γ and a guarded LTL-FO formula $\varphi_0 = \forall \bar{x} \psi_0(\bar{x})$ over the schema of Γ . To check that every run of Γ satisfies φ_0 we equivalently verify that there is no run of Γ satisfying $\neg \varphi_0 = \exists \bar{x} \neg \psi_0(\bar{x})$. Let \bar{c} be a tuple of constants of the same arity as \bar{x} . Verifying that all runs of Γ satisfy φ_0 is equivalent to checking that no run satisfies $\psi(\bar{x} \leftarrow \bar{c})$ (the formula obtained by substituting \bar{c} for \bar{x} in $\psi(\bar{x})$) for any choice of constants \bar{c} . Let us denote $\psi(\bar{x} \leftarrow \bar{c})$ by $\psi_{\bar{c}}$. Consider now a maximal subformula ξ of $\psi_{\bar{c}}$ that contains no temporal operator, which we call an FO component of $\psi_{\bar{c}}$. Note that ξ has no free variables (as variables previously free in ξ have been replaced by the constants in \bar{c}). Thus, ξ can be evaluated to true or false in every configuration of a run of Γ . This allows treating every such ξ as a proposition. More precisely, for each FO component ξ of $\psi_{\bar{c}}$, let p_ξ be a propositional symbol. Let $\psi_{\bar{c}}^{aux}$ be the LTL formula obtained by replacing in $\psi_{\bar{c}}$ every FO component ξ by p_ξ . For each configuration of Γ , the truth value of p_ξ is defined as the truth value of ξ . Clearly, a run of Γ satisfies $\psi_{\bar{c}}$ iff it satisfies $\psi_{\bar{c}}^{aux}$. Specifically, for $i \geq 0$, let $\sigma(\rho_i)$ be the truth assignment to the propositions in $\psi_{\bar{c}}^{aux}$ such that p_ξ is true iff $\rho_i \models \xi$, and let $\sigma(\rho) = \{\sigma(\rho_i)\}_{i \geq 0}$. Let us denote by $A_{\psi_{\bar{c}}}$ the Büchi automaton corresponding to the propositional LTL formula $\psi_{\bar{c}}^{aux}$. A run of $A_{\psi_{\bar{c}}}$ on $\sigma(\rho)$

is an infinite sequence of states $q_0, s_0, s_1, \dots, s_i, \dots$ such that q_0 is the start state of $A_{\psi_{\bar{c}}}$, and $\langle q_0, \sigma(\rho_0), s_0 \rangle, \langle s_i, \sigma(\rho_{i+1}), s_{i+1} \rangle$ are transitions in $A_{\psi_{\bar{c}}}$ for each $i \geq 0$. Clearly, $\rho \models \psi_{\bar{c}}$ iff there exists a run of $A_{\psi_{\bar{c}}}$ on input $\sigma(\rho)$ that goes through some accepting state, say f , infinitely often.

The core difficulty of verifying the above is that Γ is an infinite-state system (as it has infinitely many configurations), rather than a finite-state system as in classical model checking. Thus, exhaustive exploration of all possible runs of Γ is impossible. The solution lies in avoiding explicit exploration of the state space. Instead of materializing a full initial database and exploring the possible runs on it, we generate symbolic representations of equivalence classes of actual runs, called *pseudoruns*, in which every configuration retains just the information needed to check satisfaction of $\psi_{\bar{c}}$. Moreover, this information is sufficient to obtain the same information about the *next* configuration (so it is a “closed” representation system with respect to transitions of Γ). This makes crucial use of the guarded restriction. Furthermore, it can be shown that each pseudorun satisfying $\psi_{\bar{c}}$ corresponds to an actual run that also satisfies $\psi_{\bar{c}}$, on *some* possibly very large database, which is however never explicitly constructed. Thus, it is enough to limit ourselves to exploration of pseudoruns. Since pseudorun configurations turn out to be of polynomial size, this yields a PSPACE verification algorithm. We next outline the technical development in more detail.

Let $\Gamma = \langle DB, R, S, \pi, \psi, \mathcal{S} \rangle$ be a guarded singleton artifact system, and φ a guarded LTL-FO formula. Let $\psi_{\bar{c}}$ be obtained from φ as described above, for some sequence \bar{c} of constants. Let C be the set of all constants used in the specification of Γ and in $\psi_{\bar{c}}$ (this includes \bar{c}). For the purpose of this proof, we define the *active domain* of every database instance to consist of C together with the elements occurring in the instance.

For technical reasons, we will use a slightly different notion of run of Γ than in Section 2, called C -run, that depends on C and includes some additional information.

Specifically, a C -run of Γ on database DB is a sequence $\rho = \{ \langle DB, l_i, O_i, S_i \rangle \}_{i \geq 0}$ where

- $\{ \langle DB, l_i, S_i \rangle \}_{i \geq 0}$ is a run of Γ , and
- $O_i = l_{i+1}$ for $i \geq 0$.

For a run ρ , we denote its i -th configuration $\langle DB, l_i, O_i, S_i \rangle$ by ρ_i .

Intuitively, each configuration in a C -run includes both the input artifact consumed and the output artifact produced at each step. Of course, the output artifact at step i coincides with the input artifact at step $i + 1$. Recall that both the pre-condition for l_i and the post-condition for O_i are evaluated in state S_i , which makes it convenient to represent both as part of the same instance.

We say that two instances H and H' over the same database schema are C -isomorphic iff there exists an isomorphism from H to H' that is the identity on C and preserves \leq . The C -isomorphism type of H consists of all instances H' that are C -isomorphic to H . The *restriction* of an instance H to a set M of domain elements, denoted $K|_M$, is the instance consisting of the tuples in K using only elements in M .

Let $DB_i = DB|_{\text{adom}(I_i \cup O_i)}$ and $\leq_i = \leq|_{\text{adom}(I_i \cup O_i)}$. A key observation towards building pseudoruns is that, due to the guarded restriction, the truth value of each FO component of $\psi_{\bar{c}}$ in a configuration $\langle DB, l_i, O_i, S_i \rangle$ is completely determined by the isomorphism type of $\langle DB_i, l_i, O_i, S_i | C, \leq_i \rangle$. Since l_i and O_i contain only one tuple each, the number of such C -isomorphism types is finite. Moreover, due to the guarded restriction on the state rules, the same information about the possible $\langle DB_{i+1}, l_{i+1}, O_{i+1}, S_{i+1} \rangle$ is determined by the corresponding information about $\langle DB_i, l_i, O_i, S_i \rangle$. With some care, similar observations extend to satisfaction of π and ψ . Indeed, let $\zeta = \pi \wedge \psi$. Recall that ζ is existentially quantified. By a simple syntactic manipulation, we can additionally enforce that all existential quantifications in ζ are guarded. This is done by first extending R in order to include columns for all bound variables of ζ , then projecting out these columns by a guarded quantification whenever R is used in any formula. Consequently, we can assume that I_i and O_i include all witnesses needed to satisfy ζ on $\langle DB, l_i, S_i \rangle$. Thus, it is enough to consider C -isomorphism types as above. Our pseudoruns, defined shortly, will essentially represent such C -isomorphism types.

For each $\rho_i = \langle \text{DB}, l_i, O_i, S_i \rangle$, let $\rho_i^\downarrow = \langle \text{DB}_i, l_i, O_i, S_i | C, \leq_i \rangle$. We refer to the sequence $\{\rho_i^\downarrow\}_{i \geq 0}$ as the *local C-run*⁴ of ρ . We say that a sequence $\{\rho_i^\downarrow\}_{i \geq 0}$ is a local *C-run* of Γ on DB if it is the local *C-run* of some *C-run* of Γ on DB. We will use the following.

Lemma B.2 *Let Γ be a guarded singleton artifact system, φ a guarded LTL-FO formula, and $\bar{c}, \psi_{\bar{c}}, C$, and ρ be as above. Let ξ be an FO component of $\psi_{\bar{c}}$. Then for each configuration ρ_i in the *C-run* ρ , $\rho_i \models \xi$ iff $\rho_i^\downarrow \models \xi$.*

Proof: Let $\rho_i = \langle \text{DB}, l_i, O_i, S_i \rangle$. We show by induction the following:

- (\dagger) for every subformula $\xi'(\bar{x})$ of ξ with free variables \bar{x} , and sequence \bar{e} of elements in C^i of the same arity as \bar{x} , $\rho_i \models \xi'(\bar{x} \leftarrow \bar{e})$ iff $\rho_i^\downarrow \models \xi'(\bar{x} \leftarrow \bar{e})$.

As a consequence of (\dagger), $\rho_i \models \xi$ iff $\rho_i^\downarrow \models \xi$, since ξ has no free variables.

Consider (\dagger). We can assume wlog that ξ uses only \wedge, \neg and \exists . For the basis, suppose $\xi'(\bar{x})$ is an atom $Q(t_1, \dots, t_m)$ where each t_i is an element in C or a variable in \bar{x} . If Q is the state relation S , all t_i 's are elements in C by the guarded restriction, so (\dagger) holds because ρ_i^\downarrow retains $S_i | C$. If Q is R or a database relation, then again (\dagger) holds because ρ_i^\downarrow retains l_i and DB_i . The case when $\xi'(\bar{x})$ is an atom of the form $t_i = t_j$ or $t_i \leq t_j$ is obvious. Consider the induction step. If $\xi' = \xi_1 \wedge \xi_2$ or $\xi' = \neg \xi_1$ and ξ_1, ξ_2 satisfy (\dagger), it immediately follows that ξ' satisfies (\dagger). Now suppose $\xi'(\bar{x}) = \exists y (R(t_1, \dots, t_k) \wedge \varphi(\bar{x}, y))$ where each variable among the t_i 's is either y or in \bar{x} (at least one t_i is y by the guarded restriction), and (\dagger) holds for $\varphi(\bar{x}, y)$. Then $\rho_i \models \xi'(\bar{x} \leftarrow \bar{e})$ iff there exists c occurring in R such that $\rho_i \models R(t_1, \dots, t_k)[\bar{x} \leftarrow \bar{e}, y \leftarrow c] \wedge \varphi(\bar{x}, y)[\bar{x} \leftarrow \bar{e}, y \leftarrow c]$. By the induction hypothesis, this happens iff $\rho_i^\downarrow \models R(t_1, \dots, t_k)[\bar{x} \leftarrow \bar{e}, y \leftarrow c] \wedge \varphi(\bar{x}, y)[\bar{x} \leftarrow \bar{e}, y \leftarrow c]$, so $\rho_i^\downarrow \models \xi'(\bar{x} \leftarrow \bar{e})$, which shows (\dagger). \square

Lemma B.2 shows that, as desired, in a configuration ρ_i , the information relevant to satisfaction of $\psi_{\bar{c}}$ is captured by ρ_i^\downarrow . Consequently, a *C-run* satisfies $\psi_{\bar{c}}$ iff its local *C-run* satisfies $\psi_{\bar{c}}$. We next show that one can further focus on special kinds of local *C-runs*, that we call *db-periodic*. The key intuition behind db-periodicity is that in every infinite run on a finite database, the portions of local configurations that use only elements from the database must repeat at some point. However, the run may generally also involve infinitely many domain elements not occurring in the database, so proper periodicity may not be achievable. This leads to the following.

Definition B.3 *Let Γ be a guarded singleton artifact system, φ a guarded LTL-FO formula, and C be as above. A local *C-run* $\rho^\downarrow = \{\rho_i^\downarrow\}_{i \geq 0}$ of Γ on database DB is *db-periodic* if there exist $n, p > 0$ such that for every $h \geq n$, $\text{adom}(\rho_h^\downarrow) \cap \text{adom}(\text{DB}) = \text{adom}(\rho_{h+p}^\downarrow) \cap \text{adom}(\text{DB})$ and there exists a *C-isomorphism* from ρ_h^\downarrow to ρ_{h+p}^\downarrow that is the identity on $\text{adom}(\rho_h^\downarrow) \cap \text{adom}(\rho_{h+p}^\downarrow)$.*

We can show the following.

Lemma B.4 *Let Γ be a guarded singleton artifact system, φ a guarded LTL-FO formula, and $\bar{c}, \psi_{\bar{c}}$, and C be as above. There exists a *C-run* of Γ satisfying $\psi_{\bar{c}}$ iff there exists a *db-periodic local C-run* of Γ satisfying $\psi_{\bar{c}}$.*

Proof: By Lemma B.2, there exists a *C-run* of Γ satisfying $\psi_{\bar{c}}$ iff there exists a local *C-run* of Γ satisfying $\psi_{\bar{c}}$. It remains to show that if there exists a local *C-run* of Γ satisfying $\psi_{\bar{c}}$ then there exists a *db-periodic local C-run* of Γ satisfying $\psi_{\bar{c}}$. Consider a *C-run* $\rho = \{\rho_i\}_{i \geq 0}$ of Γ on database DB, satisfying $\psi_{\bar{c}}$, and the corresponding local *C-run* $\rho^\downarrow = \{\rho_i^\downarrow\}_{i \geq 0}$. Let $A_{\psi_{\bar{c}}}$ be the Büchi automaton corresponding to the propositional LTL formula $\psi_{\bar{c}}^{\text{aux}}$. Recall that for $i \geq 0$, $\sigma(\rho_i^\downarrow)$ denotes the truth assignment to the propositions in $\psi_{\bar{c}}^{\text{aux}}$ such that p_ξ is true iff $\rho_i^\downarrow \models$

⁴Our local run is an extension of the notion of local run introduced in [31, 32] for ASM transducers.

ξ , and $\sigma(\rho^\downarrow) = \{\sigma(\rho_i^\downarrow)\}_{i \geq 0}$. Since $\rho^\downarrow \models \psi_{\bar{c}}$, there exists a C -run $q_0, s_0, s_1, \dots, s_i, \dots$ of $A_{\psi_{\bar{c}}}$ on input $\sigma(\rho^\downarrow)$ that goes through some accepting state, say f , infinitely often. We define the following equivalence relation on D : $a \equiv b$ iff for every $e \in \text{adom}(\text{DB})$, $a \leq e$ iff $b \leq e$. Clearly, \equiv has finitely many equivalence classes for given C and DB . Using this, and the fact that $\text{adom}(\text{DB})$ is finite, it is easily seen by a pigeon-hole argument that there must exist $n < m$, such that $s_n = s_m = f$, $\text{adom}(\rho_n^\downarrow) \cap \text{adom}(\text{DB}) = \text{adom}(\rho_m^\downarrow) \cap \text{adom}(\text{DB})$, and ρ_n^\downarrow and ρ_m^\downarrow are isomorphic by an isomorphism α that is the identity on $\text{adom}(\rho_n^\downarrow) \cap \text{adom}(\rho_m^\downarrow)$ and for which $a \equiv \alpha(a)$ for every a . Since \leq is a dense, countable order with no endpoints, it is easily seen by a simple back-and-forth construction that there exists a bijection $\bar{\alpha}$ on D extending α and preserving \leq . Let $p = m - n$ and consider the sequence ρ' of local configurations defined inductively as follows: $\rho'_j = \rho_j^\downarrow$ for $0 \leq j \leq n + p$, and $\rho'_j = \bar{\alpha}(\rho'_{j-p})$ for $j > n + p$. It is easily seen that ρ' is a db-periodic local C -run of Γ satisfying $\psi_{\bar{c}}$. \square

We next develop a symbolic representation for local C -runs, leading to our notion of pseudorun. The idea is to represent local configurations using a fixed, finite set of symbols. Intuitively, symbols must be “reused” in such a representation, so the same symbol occurring in different symbolic configurations may correspond to different domain elements in a real local C -run.

Let $k = 2 \cdot \text{arity}(R)$. Let $V_k = C \cup \{v_1, \dots, v_k\}$, where v_1, \dots, v_k are distinct new symbols. We can clearly represent the C -isomorphism type of ρ_i^\downarrow by an instance whose domain is V_k . To do so, it is enough to fix some injective mapping f_i from $\text{adom}(\rho_i^\downarrow)$ to V_k that fixes C , and consider the instance $\tau_i = f_i(\rho_i^\downarrow)$ over V_k . By definition, τ_i is C -isomorphic to ρ_i^\downarrow . We wish to extend this to a representation of entire local C -runs by sequences of instances with domain V_k . To do so, we proceed as follows. Let $\{\rho_i^\downarrow\}_{i \geq 0}$ be a local C -run, where $\rho_i^\downarrow = \langle \text{DB}_i^p, \text{I}_i^p, \text{O}_i^p, \text{S}_i^p, \preceq_i \rangle$. We define by induction a sequence of one-to-one mappings $\{f_i\}_{i \geq 0}$, where f_i maps $\text{adom}(\rho_i^\downarrow)$ to V_k and is the identity on C :

- f_0 is an arbitrary one-to-one mapping from $\text{adom}(\rho_0^\downarrow)$ to V_k that fixes C .
- for $i \geq 0$, f_{i+1} is an arbitrary extension of $f_i|_{\text{adom}(\text{O}_i)}$ to a one-to-one mapping from $\text{adom}(\rho_{i+1}^\downarrow)$ to V_k that is the identity on C .

Now let $\tau_i = f_i(\rho_i^\downarrow)$ for each $i \geq 0$. By definition, τ_i and ρ_i^\downarrow are C -isomorphic, and the sequence $\tau = \{\tau_i\}_{i \geq 0}$ uses only elements in V_k . Also note that, since τ_i and ρ_i^\downarrow are C -isomorphic, τ satisfies $\psi_{\bar{c}}$ iff ρ^\downarrow satisfies $\psi_{\bar{c}}$.

We would like to be able to generate sequences $\{\tau_i\}_{i \geq 0}$ of instances using elements in V_k independently of any particular local C -run. To this end, we need to understand which sequences of symbolic configurations correspond to local C -runs. As a first step, we introduce the notion of C -pseudorun.

Definition B.5 Let $\Gamma = \langle \text{DB}, R, S, \pi, \psi, S \rangle$ be a guarded singleton artifact system, and let $\psi_{\bar{c}}$, C , and V_k be defined as above. A C -pseudorun of Γ is a sequence of instances

$$\{\langle \text{DB}_i, \text{I}_i, \text{O}_i, \text{S}_i, \preceq_i \rangle\}_{i \geq 0}$$

with elements in V_k such that for each $i \geq 0$:

1. $\text{DB}_i, \text{I}_i, \text{O}_i, \text{S}_i$ are database, input, output, and state instances;
2. $\text{adom}(\text{S}_i) \subseteq C$;
3. I_i and O_i contain one tuple each, and $\text{adom}(\text{DB}_i) \subseteq \text{adom}(\text{I}_i \cup \text{O}_i)$;
4. \preceq_i is an extension of $\leq|_C$ to $\text{adom}(\text{I}_i \cup \text{O}_i)$;
5. $\text{DB}_i|_{\text{adom}(\text{O}_i)} = \text{DB}_{i+1}|_{\text{adom}(\text{O}_i)}$, and $\preceq_i|_{\text{adom}(\text{O}_i)} = \preceq_{i+1}|_{\text{adom}(\text{O}_i)}$;

6. $\langle \text{DB}_i, l_i, S_i, \preceq_i \rangle \models \pi$ (with \leq interpreted as \preceq_i);
7. if $O_i = \{\langle \bar{a} \rangle\}$, then $\langle \text{DB}_i, l_i, S_i, \preceq_i \rangle \models \psi(\bar{x} \leftarrow \bar{a})$ (with \leq interpreted as \preceq_i);
8. $S_0 = \emptyset$;
9. $S_{i+1} = S'_{i+1}|_C$, where S'_{i+1} is the state relation defined from the instance $\langle \text{DB}_i, l_i, S_i, \preceq_i \rangle$, according to the state rules of Γ (with active domain semantics and \leq interpreted as \preceq_i).

Consider a local C -run $\rho^\downarrow = \{\rho_i^\downarrow\}_{i \geq 0}$ of Γ and the corresponding sequence $\{f_i(\rho_i^\downarrow)\}_{i \geq 0}$ of configurations over V_k defined above. We next show that, as desired, this is a C -pseudorun of Γ .

Lemma B.6 *Let $\rho^\downarrow = \{\rho_i^\downarrow\}_{i \geq 0}$ be a local C -run of Γ . The sequence $\{f_i(\rho_i^\downarrow)\}_{i \geq 0}$ defined above is a C -pseudorun of Γ .*

Proof: Let $\tau_i = f_i(\rho_i^\downarrow)$, $i \geq 0$. Parts (1-8) of Definition B.5 are obviously satisfied. Consider (9). Consider τ_i and τ_{i+1} for $i \geq 0$. Suppose $\phi^+(\bar{x})$ and $\phi^-(\bar{x})$ are the guarded formulas of Γ defining the tuples to be inserted, respectively deleted from S . Let \bar{e} be a sequence of elements in C of the same arity as \bar{x} . Since ρ_i^\downarrow is C -isomorphic to τ_i and ϕ^+, ϕ^- are input bounded, one can show similarly to (\dagger) in the proof of Lemma B.2 that $\tau_i \models \phi^+(\bar{e})$ iff $\rho_i^\downarrow \models \phi^+(\bar{e})$, and also $\tau_i \models \phi^-(\bar{e})$ iff $\rho_i^\downarrow \models \phi^-(\bar{e})$. Also by (\dagger) , $\rho_i^\downarrow \models \phi^+(\bar{e})$ iff $\rho_i \models \phi^+(\bar{e})$ and $\rho_i^\downarrow \models \phi^-(\bar{e})$ iff $\rho_i \models \phi^-(\bar{e})$. It follows that \bar{e} is inserted/deleted from S in the transition from ρ_i to ρ_{i+1} iff it is inserted/deleted in the transition from ρ_i^\downarrow to ρ_{i+1}^\downarrow iff it is inserted/deleted in the transition from τ_i to τ_{i+1} according to the state rule. Since by definition S is the same in $\rho_i, \rho_i^\downarrow$, and τ_i , and S is also the same in $\rho_{i+1}, \rho_{i+1}^\downarrow$, and τ_{i+1} , (9) holds. \square

Consider a C -pseudorun τ and the elements in V_k . Some of the elements occurring in different configurations of τ represent the same value in every local C -run corresponding to it, while others are independent of each other. We denote by $\langle i, a \rangle$ the occurrence of a in τ_i , where $i \geq 0$ and $a \in V_k$. Let $V_k^i = \{\langle i, a \rangle \mid a \in V_k\}$, and $V_k^\infty = \cup_{i \geq 0} V_k^i$. To capture the required equalities among elements in different configurations, we define the following equivalence relation \approx on V_k^∞ . First, for $a \in V_k - C$, let $\langle i, a \rangle \sim \langle i+1, a \rangle$ if a occurs in O_i . For $a \in C$, let $\langle i, a \rangle \sim \langle j, a \rangle$ for all i, j . Next, let \approx be the symmetric, reflexive, transitive closure of \sim . The *span* of $\langle i, a \rangle$ is the set $\text{span}(\langle i, a \rangle) = \{j \mid \langle i, a \rangle \approx \langle j, a \rangle\}$. Note that for each $a \in \text{adom}(\tau_i)$, $\text{span}(\langle i, a \rangle)$ is an interval containing i , possibly infinite to the right. We denote the equivalence class of $\langle i, a \rangle$ with respect to \approx by $\langle i, a \rangle_\approx$.

Let $\tau = \{\tau_i\}_{i \geq 0}$ be a C -pseudorun. For each $i \geq 0$, we denote by $\tau_{i, \approx}$ the result of replacing each a in τ_i by $\langle i, a \rangle_\approx$, and by τ_\approx the sequence $\{\tau_{i, \approx}\}_{i \geq 0}$.

Let \preceq_i^∞ be the total order relation on V_k^i / \approx defined by $\langle i, a \rangle_\approx \preceq_i^\infty \langle i, b \rangle_\approx$ iff $a \preceq_i b$. Let \preceq^∞ be the transitive closure of $\cup_{i \geq 0} \preceq_i^\infty$.

We note the following.

Lemma B.7 *Let τ be a C -pseudorun. Then \preceq^∞ is a partial order on V_k^∞ / \approx .*

Proof: To show that \preceq^∞ is a partial order, it is enough to show that $\cup_{i \geq 0} \prec_i^\infty$ is acyclic. Suppose, to the contrary, that there exist $a_{i_j} \in \text{adom}(\tau_{i_j})$, $1 \leq j \leq n$, such that

$$\langle i_1, a_{i_1} \rangle_\approx \prec_{i_1}^\infty \langle i_2, a_{i_2} \rangle_\approx \prec_{i_2}^\infty \langle i_3, a_{i_3} \rangle_\approx \dots \prec_{i_{n-1}}^\infty \langle i_n, a_{i_n} \rangle_\approx \prec_{i_n}^\infty \langle i_1, a_{i_1} \rangle_\approx$$

and n is minimum with this property. We can assume, wlog, that $\min(\text{span}(\langle i_1, a_{i_1} \rangle_\approx)) \leq \min(\text{span}(\langle i_j, a_{i_j} \rangle_\approx))$ for all j , $1 \leq j \leq n$. We further show that

(\dagger) $\min(\text{span}(\langle i_j, a_{i_j} \rangle)) \leq \min(\text{span}(\langle i_{j+1}, a_{i_{j+1}} \rangle))$ for all $1 \leq j < n$.

We use the minimality of n . Indeed, suppose (\dagger) does not hold, and let j be minimum such that

$$\min(\text{span}(\langle i_j, a_{i_j} \rangle)) > \min(\text{span}(\langle i_{j+1}, a_{i_{j+1}} \rangle)).$$

By assumption, $j > 1$ and by minimality of j , $\min(\text{span}(\langle i_{j-1}, a_{i_{j-1}} \rangle)) \leq \min(\text{span}(\langle i_j, a_{i_j} \rangle))$. Since $\text{span}(\langle i_j, a_{i_j} \rangle) \cap \text{span}(\langle i_{j-1}, a_{i_{j-1}} \rangle) \neq \emptyset$ and $\text{span}(\langle i_j, a_{i_j} \rangle) \cap \text{span}(\langle i_{j+1}, a_{i_{j+1}} \rangle) \neq \emptyset$ it follows that

$$m = \min(\text{span}(\langle i_j, a_{i_j} \rangle)) \in \text{span}(\langle i_{j-1}, a_{i_{j-1}} \rangle) \cap \text{span}(\langle i_{j+1}, a_{i_{j+1}} \rangle).$$

Thus, $\langle i_{j-1}, a_{i_{j-1}} \rangle_{\approx} \prec_m^{\infty} \langle i_{j+1}, a_{i_{j+1}} \rangle_{\approx}$. This yields a shorter cycle and contradicts the minimality of n . Thus, we have shown (\dagger). From (\dagger) and the fact that $\min(\text{span}(\langle i_n, a_{i_n} \rangle)) \in \text{span}(\langle i_1, a_{i_1} \rangle)$ it follows that $\text{span}(\langle i_1, a_{i_1} \rangle) \cap \text{span}(\langle i_j, a_{i_j} \rangle) \neq \emptyset$ for every j , $1 < j \leq n$. It then follows by an easy induction that $\langle i_1, a_{i_1} \rangle_{\approx} \prec_j^{\infty} \langle i_j, a_{i_j} \rangle_{\approx}$ for all such j . In particular, $\langle i_1, a_{i_1} \rangle_{\approx} \prec_n^{\infty} \langle i_n, a_{i_n} \rangle_{\approx}$, which contradicts the assumption that $\langle i_n, a_{i_n} \rangle_{\approx} \prec_n^{\infty} \langle i_1, a_{i_1} \rangle_{\approx}$. This completes the proof. \square

We next define the analog of db-periodicity for C -pseudoruns. The intention is to characterize the C -pseudoruns that correspond to db-periodic local C -runs of Γ . To ensure the correspondence, we must require some properties that are subtle consequences of the combination of infinite, ordered domain and finite database. For example, a C -pseudorun may generally create an infinite increasing chain (wrt \preceq^{∞}) of elements of V_k^{∞} / \approx that occur in the database (called *db-bound* below). Clearly, such C -pseudoruns cannot correspond to any local C -run on a finite database, and they are ruled out by requirement (*iii*) below. Note however that infinite chains not involving db-bound elements must be allowed, since these may occur in local C -runs involving infinitely many domain values.

Definition B.8 Let Γ be a singleton artifact system and $\tau = \{\tau_i\}_{i \geq 0} = \{\langle \text{DB}_i, l_i, O_i, S_i, \preceq_i \rangle\}_{i \geq 0}$ a C -pseudorun of Γ . Consider τ_{\approx} . An equivalence class e of \approx is called *db-bound* iff there exists j such that $\langle j, a \rangle \in e$ for some a occurring in DB_j . The C -pseudorun τ is db-periodic if there exist $n, p > 0$ such that, for every $h \geq n$, there exists an isomorphism μ_h from $\tau_{h,\approx}$ to $\tau_{h+p,\approx}$ such that:

- (i) μ_h is the identity on $\text{adom}(\tau_{h,\approx}) \cap \text{adom}(\tau_{h+p,\approx})$,
- (ii) $\langle h, a \rangle_{\approx}$ is db-bound iff $\mu_h(\langle h, a \rangle_{\approx})$ is db-bound,
- (iii) if $\langle h, a \rangle_{\approx}$ is db-bound then $\langle h, a \rangle_{\approx}$ and $\mu_h(\langle h, a \rangle_{\approx})$ are incomparable wrt \prec^{∞} .

It will be useful to note the following.

Lemma B.9 Let $\tau = \{\tau_i\}_{i \geq 0} = \{\langle \text{DB}_i, l_i, O_i, S_i, \preceq_i \rangle\}_{i \geq 0}$ be a db-periodic C -pseudorun of Γ . Let n, p and μ_h be as above. Then the following hold:

- (i) for each $a \in V_k^{\infty} / \approx$ and $i, j \in \text{span}(a)$ such that $i, j \geq n$, $\mu_i(a) = \mu_j(a)$; and,
- (ii) for every $h \geq n$, $\text{adom}(\tau_{h,\approx}) \cap \text{adom}(\tau_{h+p,\approx}) = \bigcap_{i \geq n} \text{adom}(\tau_{i,\approx})$.

Proof: Both (i) and (ii) are straightforward consequences of the fact that $O_i = I_{i+1}$ for every $i \geq 0$, and $\mu_i(O_i) = O_{i+p} = I_{i+p+1} = \mu_{i+1}(I_{i+1})$ for every $i \geq n$. Thus, μ_i and μ_{i+1} agree on $O_i = I_{i+1}$. \square

We next show the following key result, establishing the desired connection between pseudoruns and actual runs.

Lemma B.10 *Let Γ be a guarded singleton artifact system and φ a guarded LTL-FO formula. Let C , \bar{c} , and $\psi_{\bar{c}}$ be as above. The following are equivalent:*

- (a) *there exists some db-periodic local C -run ρ^\downarrow of Γ such that $\rho^\downarrow \models \psi_{\bar{c}}$, and*
- (b) *there exists some db-periodic C -pseudorun τ of Γ such that $\tau \models \psi_{\bar{c}}$.*

Proof: Consider (a) \rightarrow (b). Let $\rho^\downarrow = \{\rho_i^\downarrow\}_{i \geq 0}$ be a db-periodic local C -run of Γ , that satisfies $\psi_{\bar{c}}$. Consider the C -pseudorun $\tau = \{\tau_i\}_{i \geq 0} = \{f_i(\rho_i^\downarrow)\}_{i \geq 0}$ of Γ constructed prior to Lemma B.6. Since ρ_i^\downarrow and τ_i are C -isomorphic for each $i \geq 0$, $\tau \models \psi_{\bar{c}}$. We show that τ is db-periodic. Since ρ^\downarrow is db-periodic, there exist $n, p > 0$ such that for every $h \geq n$, $\text{adom}(\rho_h^\downarrow) \cap \text{adom}(\text{DB}) = \text{adom}(\rho_{h+p}^\downarrow) \cap \text{adom}(\text{DB})$ and there exists a C -isomorphism α from ρ_h^\downarrow to ρ_{h+p}^\downarrow that is the identity on $\text{adom}(\rho_h^\downarrow) \cap \text{adom}(\rho_{h+p}^\downarrow)$. Consider the mapping g from V_k^∞ / \approx to the domain of ρ^\downarrow defined by inverting the mappings f_i , $i \geq 0$. More precisely, for each $i \geq 0$ and $a \in V_k$, let $g(\langle i, a \rangle_\approx) = f_i^{-1}(a)$. The following are immediate from the definition, for each $i, j \geq 0$:

1. g is well defined,
2. if $\langle i, a \rangle_\approx$ is db-bound then $g(\langle i, a \rangle_\approx) \in \text{adom}(\text{DB})$,
3. if $\langle i, a \rangle_\approx \prec^\infty \langle j, a \rangle_\approx$ then $g(\langle i, a \rangle_\approx) < g(\langle j, a \rangle_\approx)$.

Now consider the mapping μ which is the restriction of $g \circ \alpha \circ g^{-1}$ to $\text{adom}(\tau_{h,\approx})$. Clearly, μ is an isomorphism from $\tau_{h,\approx}$ to $\tau_{h+p,\approx}$. We show that μ satisfies the properties (i)-(iii) required by the definition of db-periodicity. Consider (i). If $\langle h, a \rangle_\approx \in \text{adom}(\tau_{h,\approx}) \cap \text{adom}(\tau_{h+p,\approx})$, then $\langle h, a \rangle_\approx \approx \langle h+p, a \rangle_\approx$ and by (1) above, $g(\langle h, a \rangle_\approx) = g(\langle h+p, a \rangle_\approx)$ so $g(\langle h, a \rangle_\approx) \in \text{adom}(\rho_h^\downarrow) \cap \text{adom}(\rho_{h+p}^\downarrow)$ so $\alpha(g(\langle h, a \rangle_\approx)) = g(\langle h, a \rangle_\approx) = g(\langle h+p, a \rangle_\approx)$ and $\mu(\langle h, a \rangle_\approx) = \langle h+p, a \rangle_\approx = \langle h, a \rangle_\approx$. Now consider (ii). Suppose $\langle h, a \rangle_\approx$ is db-bound. Then by (2), $g(\langle h, a \rangle_\approx) \in \text{adom}(\text{DB})$ so $\alpha(g(\langle h, a \rangle_\approx)) = g(\langle h, a \rangle_\approx)$ and $\mu(\langle h, a \rangle_\approx)$ is also db-bound. The converse is similar. Finally, consider (iii). Suppose $\langle h, a \rangle_\approx$ is db-bound and $\langle h, a \rangle_\approx \prec^\infty \mu(\langle h, a \rangle_\approx)$. By (3), $g(\langle h, a \rangle_\approx) < g(\mu(\langle h, a \rangle_\approx))$. However, by definition $g(\mu(\langle h, a \rangle_\approx)) = \alpha(g(\langle h, a \rangle_\approx))$ and by (2), $g(\langle h, a \rangle_\approx) \in \text{adom}(\text{DB})$. Since α is the identity on $\text{adom}(\text{DB})$, $\alpha(g(\langle h, a \rangle_\approx)) = g(\langle h, a \rangle_\approx)$. Thus, $g(\langle h, a \rangle_\approx) < g(\langle h, a \rangle_\approx)$, which is a contradiction. The case when $\mu(\langle h, a \rangle_\approx) \prec^\infty \langle h, a \rangle_\approx$ is similar. Thus, (iii) holds, and τ is db-periodic.

Now consider the harder (b) \rightarrow (a). Let $\tau = \{\tau_i\}_{i \geq 0}$ be a db-periodic C -pseudorun satisfying $\psi_{\bar{c}}$. We define a finite database DB and a db-periodic local C -run $\rho^\downarrow = \{\rho_i^\downarrow\}_{i \geq 0}$ of Γ on DB such that for each $i \geq 0$, ρ_i^\downarrow is C -isomorphic to τ_i . In particular, $\rho^\downarrow \models \psi_{\bar{c}}$.

In order to obtain a local C -run from τ , we will assign values from D to elements in V_k^∞ / \approx . This yields a sequence of instances that obeys the definition of local C -run, except that it uses a database with an infinite domain. We will then perform some ‘‘surgery’’ on this sequence in order to obtain a proper local C -run over a finite database. The resulting local C -run may still use infinitely many domain values, but only finitely many of them occur in the database.

Let \preceq_\star^∞ be a linear extension of the partial order \preceq^∞ with the following property: for all pairs a, b of db-bound elements in V_k^∞ / \approx that are incomparable with respect to \preceq^∞ , and every $c \neq a, b$ in V_k^∞ / \approx , $c \preceq_\star^\infty a$ iff $c \preceq_\star^\infty b$ and $a \preceq_\star^\infty c$ iff $b \preceq_\star^\infty c$. It is easy to see that such \preceq_\star^∞ exists.

Next, let g be a bijection from V_k^∞ / \approx to (D, \leq) that is the identity on C and preserves \preceq_\star^∞ . Such a bijection exists because \preceq_\star^∞ is an order on a countable set and \leq is dense without endpoints. Now consider the sequence $\tau^g = \{g(\tau_{i,\approx})\}_{i \geq 0}$. We denote for conciseness $g(\tau_{i,\approx})$ by τ_i^g , and let $\tau_i^g = \langle \text{DB}_i, l_i, O_i, S_i, \leq_i \rangle$, $i \geq 0$. Let $\text{DB}^g = \cup_{i \geq 0} \text{DB}_i$. It is easily seen, using the consistency requirement (5) in the definition of pseudorun, that $\text{DB}^g | \text{adom}(\text{DB}_i) = \text{DB}_i$ for all $i \geq 0$. It follows that τ^g satisfies the definition of a local C -run of Γ on database DB^g , except for the requirement that DB^g be finite.

For given τ^g , it will be convenient to extend the notion of span to elements in $\text{adom}(\tau^g)$, which is done as for pseudoruns. In particular, $\text{span}(g(a)) = \text{span}(a)$. We call an element in τ^g db-bound iff it is the image under g of a db-bound element in $\text{adom}(\tau_{\approx})$. Equivalently, a is db-bound iff $a \in \text{adom}(\text{DB}_i)$ for some $i \geq 0$. Also note that, since τ^g is isomorphic to τ_{\approx} , the analog of Lemma B.9 also holds for τ^g .

Since τ is db-periodic, there exist $n, p' > 0$ satisfying the requirements of db-periodicity of a pseudorun. Let $p = 3p'$. It is easily seen, using part (ii) of Lemma B.9, that τ_n^g and τ_{n+p}^g have the following property:

- (†) there are no $a, b, c \in \text{adom}(\tau^g)$ such that $a \in \text{adom}(\tau_n^g) - \text{adom}(\tau_{n+p}^g)$, $b \in \text{adom}(\tau_{n+p}^g) - \text{adom}(\tau_n^g)$, $c \notin \text{adom}(\tau_n^g) \cap \text{adom}(\tau_{n+p}^g)$, such that $\text{span}(a) \cap \text{span}(c) \neq \emptyset$ and $\text{span}(b) \cap \text{span}(c) \neq \emptyset$.

Intuitively, (†) states that τ_n^g and τ_{n+p}^g are sufficiently far apart that there is no interference between them (except through their common elements). We next use this to perform some “surgery” on τ^g in order to obtain a local C -run that uses only finitely many db-bound elements. This will allow us to replace the infinite DB^g by a finite database.

Let μ_n be the isomorphism from $\tau_{n,\approx}$ to $\tau_{n+p,\approx}$ whose existence is guaranteed by the db-periodicity of τ . Let μ_n^g be the C -isomorphism from τ_n^g to τ_{n+p}^g induced by μ_n via g . Let μ_{db} be the restriction of μ_n^g to the db-bound elements in $\text{adom}(\tau_n^g)$, extended to the identity everywhere else. In particular, note that μ_n^g is the identity on $\text{adom}(\tau_n^g) \cap \text{adom}(\tau_{n+p}^g)$. Consider the subsequence $\bar{\tau} = \mu_{db}(\tau_n^g), \dots, \mu_{db}(\tau_{n+p-1}^g)$. We show that

- (‡) $\mu_{db}(\tau_i^g)$ and τ_i^g are C -isomorphic by μ_{db} , $n \leq i < n + p$.

It is obvious that μ_{db} is the identity on C , since $C \subseteq \text{adom}(\tau_n^g) \cap \text{adom}(\tau_{n+p}^g)$. To establish (‡), it is enough to show that μ_{db} is injective and preserves $<$. For injectivity, suppose $\mu_{db}(a) = \mu_{db}(b)$ and $a \neq b$ for $a, b \in \text{adom}(\tau_i^g)$. Clearly, this can only happen if $a \in \text{adom}(\tau_n^g)$ and $b \notin \text{adom}(\tau_n^g)$ (or conversely, which is similar). But then $\mu_{db}(b) = b$, and so $\mu_{db}(a) = b$ and $b \in \text{adom}(\tau_{n+p}^g)$. However, this contradicts (†). Hence, μ_{db} is injective. To see that μ_{db} preserves $<$, suppose $a, b \in \text{adom}(\tau_i^g)$ and $a < b$. If a, b are db-bound and $a, b \in \text{adom}(\tau_n^g)$ then $\mu_{db}(a) = \mu_n^g(a)$ and $\mu_{db}(b) = \mu_n^g(b)$ and $\mu_n^g(a) < \mu_n^g(b)$ since μ_n^g is an isomorphism. Similarly, if a, b are not db-bound or $a, b \notin \text{adom}(\tau_n^g)$ then $\mu_{db}(a) = a$ and $\mu_{db}(b) = b$, so $\mu_{db}(a) < \mu_{db}(b)$. Suppose a is db-bound and belongs to $\text{adom}(\tau_n^g)$, and b is not db-bound or $b \notin \text{adom}(\tau_n^g)$ (the other case is symmetric). Then $\mu_{db}(b) = b$. Let μ_i^g be the isomorphism from τ_i^g to τ_{i+p}^g induced by μ_i . Since $n, i \in \text{span}(a)$, $\mu_{db}(a) = \mu_n^g(a) = \mu_i^g(a)$ by (i) of Lemma B.9. Since a is db-bound, $g^{-1}(a)$ and $\mu_i(g^{-1}(a))$ are db-bound and incomparable with respect to \prec_{\star}^{∞} . Since $a < b$, $g^{-1}(a) \prec_{\star}^{\infty} g^{-1}(b)$, and from the definition of \prec_{\star}^{∞} it follows that $\mu_i(g^{-1}(a)) \prec_{\star}^{\infty} g^{-1}(b)$. Since g preserves \prec_{\star}^{∞} , $\mu_i^g(a) < b$. Thus, $\mu_{db}(a) < \mu_{db}(b)$. This proves (‡).

Note that $\mu_{db}(\tau_n^g)$ and τ_{n+p}^g have identical restrictions to their db-bound elements (so in particular, their database and state relations are the same). We next construct a db-periodic local C -run using finitely many db-bound elements as follows. Essentially, this is done by repeating forever the sub-sequence

$$\bar{\tau} = \mu_{db}(\tau_n^g), \dots, \mu_{db}(\tau_{n+p-1}^g)$$

following $\tau_0^g, \dots, \tau_{n+p-1}^g$, except that elements that are not db-bound remain unchanged. More precisely, we define a local C -run $\rho^{\downarrow} = \{\rho_i^{\downarrow}\}_{i \geq 0}$ as follows. First, $\rho_i^{\downarrow} = \tau_i^g$ for $0 \leq i < n + p - 1$. Next, consider i such that $n + qp \leq i < n + (q + 1)p$ for $q \geq 1$. From the db-periodicity of τ and the fact that μ_{db} is a C -isomorphism, it follows that τ_i^g and $\mu_{db}(\tau_{n+(i-qp)}^g)$ are C -isomorphic by an isomorphism that is the identity on the intersection of their domains. Let ν be the restriction of this isomorphism to the db-bound elements of τ_i^g , extended with the identity on the elements that are not db-bound. Let $\rho_i^{\downarrow} = \nu(\tau_i^g)$. Thus, the db-bounded portion of τ_i^g is replaced with the db-bounded portion of $\mu_{db}(\tau_{n+(i-qp)}^g)$ while the elements that are not db-bound remain unchanged. Similarly to (‡), it can be shown that ν remains a C -isomorphism. Thus, ρ_i^{\downarrow} is C -isomorphic to τ_i^g and to τ_i for all $i \geq 0$.

Let $\rho_i^{\downarrow} = \langle \text{DB}'_i, l'_i, O'_i, S'_i, \leq'_i \rangle$, $i \geq 0$, and let $\text{DB}' = \cup_{i \geq 0} \text{DB}'_i$. Clearly, DB' is finite. We will show that ρ^{\downarrow} is a db-periodic local C -run of Γ . We claim that

(*) $DB' \upharpoonright_{\text{adom}(\rho_i^\downarrow)} = DB'_i$ for each $i \geq 0$.

This follows straightforwardly from (†) and the easily shown fact that ρ^\downarrow satisfies the consistency criterion $DB'_i \upharpoonright_{\text{adom}(\rho_i^\downarrow) \cap \text{adom}(\rho_{i+1}^\downarrow)} = DB'_{i+1} \upharpoonright_{\text{adom}(\rho_i^\downarrow) \cap \text{adom}(\rho_{i+1}^\downarrow)}$ for every $i \geq 0$.

Let $\rho = \{\langle DB', l'_i, O'_i, S''_i \rangle\}_{i \geq 0}$ be obtained from ρ^\downarrow by computing for each $i \geq 0$, S''_{i+1} from $\langle DB', l'_i, O'_i, S''_i \rangle$, using the state rules of Γ . From (*), the definition of pseudorun, and the construction of ρ^\downarrow , it is clear that ρ is a run of Γ on database DB' , and ρ^\downarrow is the local C -run of ρ . Also, ρ^\downarrow is db-periodic by construction. Since ρ_i^\downarrow is C -isomorphic to τ_i for every $i \geq 0$, and $\tau \models \psi_{\bar{c}}$, it follows that $\rho^\downarrow \models \psi_{\bar{c}}$. This completes the proof. \square

Lemma B.10 says that in order to determine whether some C -run of Γ satisfies $\psi_{\bar{c}}$, it is enough to focus on db-periodic C -pseudoruns of Γ . However, such pseudoruns are still infinite. Fortunately, as shown next, the db-periodicity guarantees the existence of certain prefixes providing finite representations for such C -pseudoruns. Thus, it is sufficient to search for such finite prefixes.

Lemma B.11 *Let Γ , C , and $\psi_{\bar{c}}$ be as above. Let $A_{\psi_{\bar{c}}}$ be the Büchi automaton corresponding to $\psi_{\bar{c}}^{aux}$. There exists a db-periodic C -pseudorun of Γ satisfying $\psi_{\bar{c}}$ iff there exist $n, p > 0$ and a finite prefix $\tau = \{\tau_i\}_{0 \leq i \leq n+p}$ of a C -pseudorun of Γ such that:*

1. *some run of $A_{\psi_{\bar{c}}}$ on $\{\sigma(\tau_i)\}_{0 \leq i \leq n+p}$ reaches an accepting state f at $\sigma(\tau_n)$ and $\sigma(\tau_{n+p})$,*
2. *there exists an isomorphism μ from $\tau_{n,\approx}$ to $\tau_{n+p,\approx}$ satisfying conditions (i) – (iii) in the definition of db-periodic C -pseudorun and additionally $\mu(\langle n, a \rangle_{\approx}) = \langle n+p, a \rangle_{\approx}$ for every $a \in V_k$, and*
3. *there is no $a \in V_k$ and db-bound $e \in \cup_{n \leq i \leq n+p} V_k^i / \approx$ such that*

$$\langle n, a \rangle_{\approx} \prec_{\tau} e \prec_{\tau} \langle n+p, a \rangle_{\approx} \text{ or } \langle n+p, a \rangle_{\approx} \prec_{\tau} e \prec_{\tau} \langle n, a \rangle_{\approx}$$

where \preceq_{τ} is the partial order on $\cup_{0 \leq i \leq n+p} V_k^i / \approx$ induced by $\cup_{0 \leq i \leq n+p} \preceq_i^{\infty}$.

Proof: The *only if* part follows from a simple pigeon-hole argument on a db-periodic C -pseudorun satisfying $\psi_{\bar{c}}$. Part (3) follows from condition (iii) in the definition of db-periodic C -pseudorun. Consider the *if* part. Suppose we have a finite C -pseudorun prefix $\{\tau_i\}_{0 \leq i \leq n+p}$ satisfying (1)-(3). Note that by (2), $\tau_n = \tau_{n+p}$. We can obtain an infinite db-periodic C -pseudorun satisfying $\psi_{\bar{c}}$ by concatenating infinitely many times the subsequence $\tau_{n+1} \dots \tau_{n+p}$ following τ_n . Satisfaction by the resulting C -pseudorun of (iii) in the definition of db-periodic C -pseudorun follows from the fact that $\{\tau_i\}_{0 \leq i \leq n+p}$ satisfies (3). \square

We are now ready to describe a non-deterministic PSPACE verification algorithm for guarded singleton artifact systems and guarded LTL-FO properties. The input to the algorithm is a guarded singleton artifact system Γ and a guarded LTL-FO formula φ . Let $\exists \bar{x} \psi(\bar{x})$ be the negation of φ and let \bar{c} be a non-deterministically chosen sequence of possibly repeating constants, one for each variable in \bar{x} . Let C consist of \bar{c} together with all constants used in the specification of Γ or in φ . Let $\psi_{\bar{c}} = \psi[\bar{x} \leftarrow \bar{c}]$ and let $\psi_{\bar{c}}^{aux}$ be the propositional LTL formula obtained by replacing each FO component ξ of $\psi_{\bar{c}}$ by a propositional symbol p_{ξ} . Let $A_{\psi_{\bar{c}}}$ be the Büchi automaton corresponding to $\psi_{\bar{c}}^{aux}$. Let $k = 2 \cdot \text{arity}(R)$ and $V_k = C \cup \{v_1, \dots, v_k\}$, where the v_i 's are distinct new constants. We will use a non-deterministic PSPACE algorithm to guess a finite C -pseudorun prefix as in Lemma B.11. This will be done by non-deterministically generating consecutive configurations in the pseudorun while running $A_{\psi_{\bar{c}}}$ on the generated configurations. In order to check that properties (2) and (3) are satisfied, we will need to incrementally maintain some additional information. Specifically, let π be a finite sequence $\pi_1 \dots \pi_m$ of consecutive pseudoconfigurations. Similarly to Lemma B.11, let \preceq_{π} be the partial order on $\cup_{1 \leq i \leq m} V_k^i / \approx$ induced by $\cup_{1 \leq i \leq m} \preceq_i^{\infty}$. We define four partial mappings from the elements of π_1 to those of π_m :

- min_π is the mapping from V_k to V_k such that $min_\pi(a) = b$ iff $\langle m, b \rangle_\approx$ is the minimum element in V_k^m / \approx for which $\langle 1, a \rangle_\approx \preceq_\pi \langle m, b \rangle_\approx$
- min_π^{db} is the mapping from V_k to V_k such that $min_\pi(a) = b$ iff $\langle m, b \rangle_\approx$ is the minimum element in V_k^m / \approx for which exists a db-bound $e \in \cup_{1 \leq i \leq m} V_k^i / \approx$ such that $\langle 1, a \rangle_\approx \prec_\pi e \prec_\pi \langle m, b \rangle_\approx$
- max_π is the mapping from V_k to V_k such that $max_\pi(a) = b$ iff $\langle m, b \rangle_\approx$ is the maximum element in V_k^m / \approx for which $\langle m, b \rangle_\approx \succeq_\pi \langle 1, a \rangle_\approx$
- max_π^{db} is the mapping from V_k to V_k such that $max_\pi(a) = b$ iff $\langle m, b \rangle_\approx$ is the maximum element in V_k^m / \approx for which there exists a db-bound $e \in \cup_{1 \leq i \leq m} V_k^i / \approx$ such that $\langle m, b \rangle_\approx \prec_\pi e \prec_\pi \langle 1, a \rangle_\approx$.

Clearly, $min_\pi(a) \preceq_\pi min_\pi^{db}(a)$ and $max_\pi^{db}(a) \preceq_\pi max_\pi(a)$ for every $a \in V_k$. It is easily seen that the four mappings can be maintained incrementally, in the following sense. For $f \in \{min, min^{db}, max, max^{db}\}$, a finite sequence π of consecutive pseudoconfigurations, and a pseudoconfiguration δ consecutive to the last configuration in π , $f_{\pi\delta}$ can be computed from f_π and δ in polynomial time. In addition to the four mappings, we need to maintain the sets of db-bound elements in π_1 and π_m . Specifically, let $B_1 = \{c \in V_k \mid \langle 1, c \rangle_\approx \text{ is db-bound in } \pi\}$, and $B_2 = \{c \in V_k \mid \langle m, c \rangle_\approx \text{ is db-bound in } \pi\}$. Note that both sets may change when a new pseudoconfiguration is appended to π . It is easily seen that B_1 and B_2 can also be incrementally maintained.

We use the following non-deterministic PSPACE algorithms:

- Büchi-Next: on input $(\psi_{\bar{c}}^{aux}, s, \sigma)$, where s is a state of $A_{\psi_{\bar{c}}}$ and σ is a truth assignment to the propositions in $\psi_{\bar{c}}^{aux}$, the algorithm returns a state s' of $A_{\psi_{\bar{c}}}$ such that $\langle s, \sigma, s' \rangle$ is a transition in $A_{\psi_{\bar{c}}}$.
- Pseudorun-Next: given as input a configuration τ in a C -pseudorun of Γ , output a possible next configuration τ' in the pseudorun.

The algorithm now proceeds as follows:

1. flag := 0;
2. set τ_0 to an initial configuration of a C -pseudorun of Γ ;
3. initialize each of the mappings $f \in \{min, min^{db}, max, max^{db}\}$ to f_{τ_0} , and B_1, B_2 to the set of db-bound elements in τ_0 ;
4. set s_0 to some output of Büchi-Next($\psi_{\bar{c}}, q_0, \sigma(\tau_0)$), where q_0 is the start state of $A_{\psi_{\bar{c}}}$;
5. set (s, τ) to (s_0, τ_0) ;
6. if flag = 0 and s is an accepting state of $A_{\psi_{\bar{c}}}$ then non-deterministically continue or set $(\bar{s}, \bar{\tau})$ to (s, τ) and set flag:= 1;
7. set τ to Pseudorun-Next(τ) and update $min, min^{db}, max, max^{db}, B_1$ and B_2 ;
8. set s to Büchi-Next($\psi_{\bar{c}}, s, \sigma(\tau)$);
9. if flag = 1, $(s, \tau) = (\bar{s}, \bar{\tau})$, and the following conditions are satisfied:
 - $B_1 = B_2$ (let $B = B_1 = B_2$);
 - for each $a \in B$, $min(a)$ is undefined or $a \prec_\tau min(a)$, and $max(a)$ is undefined or $max(a) \prec_\tau a$,

- for each $a \in V_k - B$, $\text{min}^{db}(a)$ is undefined or $a \prec_\tau \text{min}^{db}(a)$, and $\text{max}^{db}(a)$ is undefined or $\text{max}^{db}(a) \prec_\tau a$;

output YES and stop. Otherwise, go to 6.

Clearly, the above nondeterministic PSPACE algorithm outputs YES iff there exists a finite prefix of a C -pseudorun of Γ satisfying the conditions of Lemma B.11. This in turn happens iff there exists a db-periodic C -pseudorun of Γ accepted by $A_{\psi_{\bar{c}}}$. Finally, in view of Lemmas B.10 and B.4, this holds iff there exists a run of Γ satisfying $\psi_{\bar{c}}$.

Observe that if the arity of relations in the schema of Γ is not bounded, the above algorithm is in EXSPACE. This establishes the upper bound in Theorem 3.3. The PSPACE-hardness (in the case of fixed arities) is a straightforward reduction from Quantified Boolean Formula [14], similar to the proof in [32] (details are omitted).

C Proofs of Section 4

Proof of Theorem 4.1 The proof is by reduction from the Post Correspondence Problem (PCP), known to be undecidable [29]. Consider an instance of the PCP, i.e. two sequences of non-empty words u_1, \dots, u_k and v_1, \dots, v_k over alphabet $\{0, 1\}$, for $k > 0$. Recall that a solution of the PCP is a sequence of indexes $i_1, \dots, i_n \in [1, k]$ such that $u_{i_1} \dots u_{i_n} = v_{i_1} \dots v_{i_n}$. We construct a guarded, singleton, hybrid-attribute artifact system Γ with one relational attribute, and an LTL-FO formula φ , such that $\Gamma \models \varphi$ iff there is no solution to the PCP instance. Adapting our notation for singleton artifact systems to hybrid-attribute systems, we denote Γ by $\langle \mathcal{DB}, R, S, \pi, \psi_R, \psi_S \rangle$ (there are now separate post-conditions ψ_R and ψ_S for R and S , and no state rules). We now describe Γ informally. The database \mathcal{DB} represents a directed graph whose nodes are labeled by 0, 1. To this end, \mathcal{DB} consists of a binary relation E for the edges, and a unary relation One representing the labeling ($One(v)$ means v is labeled 1 and $\neg One(v)$ means v is labeled 0). The relational attribute set S has a single attribute (so S holds a set). Intuitively, the role of S is to help identify a simple path within E . Once this is done, the labels along the path determine a word w over $\{0, 1\}$. Then it remains to guess a parsing of the word as $u_{i_1} \dots u_{i_n}$ and simultaneously as $v_{i_1} \dots v_{i_n}$ for some $n > 0$ and $i_j \in [1, k], 1 \leq j \leq n$. We now provide more details.

The tuple attribute set R consists of six attributes: $Stage, Last, Last', Index, L_u, L_v$. Intuitively, $Stage$ is used to identify different stages in the simulation, so essentially serves as a finite-state non-deterministic control. $Last$ and $Last'$ are used in the construction of the path, and $Index, L_u$ and L_v are used in the matching phase. We begin with the construction of the simple path in E , beginning at a designated node $\mathbf{start} \in D$. Edges are added to the path non-deterministically, one at a time. The most recently added node is held in the attribute $Last$, initialized to \mathbf{start} . Attribute $Last'$ lags one step behind $Last$ (so holds the next-to-last node added to the path). The set S is initialized to contain \mathbf{start} and all nodes x for which $\langle x, \mathbf{start} \rangle \in E$. $Last, Last'$ and S are then updated inductively as follows. The update occurs in alternation, first for $Last$ and $Last'$, and then for S (this is controlled by attribute $Stage$). If $Last = y$, then the value of $Last$ in the output is an arbitrary $z \neq y$ such that $\langle y, z \rangle \in E$ and $z \notin S$ (and the value of $Last'$ becomes y). Next, S is updated by adding to it all nodes $x \neq y$ for which $\langle x, z \rangle \in E$, as well as all nodes $x \neq z$ for which $\langle y, x \rangle \in E$ (z and y are available in $Last$ and $Last'$, respectively). This is repeated until it is ended non-deterministically (using attribute $Stage$). Note that at the end of this stage, S uniquely determines a simple path P in E from \mathbf{start} to the node in $Last$.

The next stage consists of matching P to a solution of the PCP. We use two attributes L_u, L_v of R to hold, respectively, the last nodes of P reached by u -words and v -words (both L_u and L_v are initialized to \mathbf{start}). These are updated inductively as follows. First, we guess an index $i \in [1, k]$ using the attribute $Index$ of R . Then we advance L_u along P by $|u_i|$ nodes so that the labels spell u_i , and L_v by $|v_i|$ nodes along P so that the labels spell v_i (this is specified in the post-condition for R). If at any point L_u and L_v contain the same node

$e \neq \mathbf{start}$, then a successful matching has been found. If this happens in a configuration, the configuration is repeated forever. Note that otherwise all partial runs of Γ block after finitely many steps.

Clearly, the PCP has a solution iff there exists an (infinite) run of Γ . Let φ be the LTL-FO property **false**. Thus, the PCP has no solution iff $\Gamma \models \mathbf{false}$. □