# DB-Nets: on The Marriage of Colored Petri Nets and Relational Databases

Marco Montali and Andrey Rivkin

Free University of Bozen-Bolzano, Piazza Domenicani 3, 39100 Bolzano, Italy montali,rivkin@inf.unibz.it

**Abstract.** The integrated management of business processes and master data is being increasingly considered as a fundamental problem, by both the academia and the industry. In this position paper, we focus on the foundations of the problem, arguing that contemporary approaches struggle to find a suitable equilibrium between data- and process-related aspects. We then propose db-nets, a new formal model that balances such two pillars through the marriage of colored Petri nets and relational databases. We invite the research community to build on this model, discussing its potential in modeling, formal verification, and simulation.

## 1 Introduction

In contemporary organizations, the integrated management of business processes (BPs) and master data is being increasingly considered as a fundamental problem, both by academia and the industry. From the practical point of view, it has been widely recognized that the traditional isolation between process and data management induces a fragmentation and redundancies in the organizational structure and its underlying IT solutions, with experts and tools solely centered around data, and others only focusing on process management [29,17,28]. This isolation falls short, especially when it comes to knowledge-intensive and humanempowered processes [19,4,21]. Foundational research has witnessed a similar trend, with completely separate fields of research either focused on the foundations of data management, or the principles underlying dynamic concurrent systems, database theory and Petri net theory being the two most prominent representatives within such fields.

To answer the increasing demand of integrated models holistically tackling the dynamics of a complex domain and the manipulation of data, a plethora of approaches has emerged in the last decade. Such approaches can be classified in two main groups, again implicitly reflecting the process-data dichotomy. A first series of approaches comes from Petri nets, the reference formalism to represent the control-flow of BPs, with the purpose of increasing their awareness of data. All such models are more or less directly inspired by a class of high-level Petri nets called Colored Petri nets (CPNs) [20,3], where colors abstractly account for data types, and where control threads (i.e., tokens) progressing through the net carry data conforming to colors. In their original formulation, though, CPNs must be severely restricted when it comes to their formal analysis, in particular requiring color domains to be finite, and in principle allowing one to get rid of the data via propositionalization. Several data-aware fragments of high-level Petri nets that are amenable to formal analysis even in the case of infinite color domains have consequently been studied, ranging from nets where tokens carry single data values (as in data- and  $\nu$ -nets [23,30]), to nets where tokens are associated to more complex data structures such as nested relations [18], nested terms [32], or XML documents [8]. The common, main issue of al such approaches is that data are still subsidiary to the control-flow dimension: data elements "locally" attached to tokens, without being connected to each other by a global, end-to-end data model. In this light, CPN-based approaches naturally support the key notion of process instance (or case) through that of token, together with the key notion of case attributes [31]. However, they do not lend themselves to modeling, querying, updating, and ultimately reasoning on persistent, relational data, like those typically maintained inside an enterprise information system. For this reason, they are unable to suitably formalize concrete BP execution engines, which all provide support for explicitly interconnecting a BP with an underlying relational persistent storage [15].

The second group of foundational approaches to data-aware processes has emerged at the intersection of database theory, formal methods and conceptual modeling, and specularly mirrors the advantages and lacks of CPN-based solutions. Such proposals go under the umbrella term of data-centric approaches [11], and gained momentum during the last decade, in particular due to the development of the business artifact paradigm [13], which also lead to concrete languages and implementations [14,21]. The common denominator of all such approaches is that processes are centered around an explicit, persistent data component maintaining information about the domain of interest, and possibly capturing also its semantics in terms of classes, relations, and constraints. Atomic tasks induce CRUD (create-ready-update-delete) operations over the data component, in turn supporting the evolution of the master data maintained therein. Proposals then differ in terms of the adopted data model (e.g., relational, tree-shaped, graph-structured), and on the nature of information (e.g., whether it is complete or not). For example, [9,16] focus on relational databases, while [7] refers to XML. The main downside of data-centric process models is that they disregard an explicit representation of how atomic tasks have to be sequenced over time, only implicitly representing the control flow via (event-)condition-action rules [16,14,9].

In this position paper, we argue that the lack of equilibrium between dataand process-related aspects in the aforementioned proposals is a major obstacle towards modeling, verifying, monitoring, mining, and ultimately *understanding* data-aware business processes. We believe that this can only be achieved by better balancing such two pillars. This, in turn, calls for the development of further modeling abstractions tailored to establishing more intimate, synergic connections between CPNs and data-centric approaches. To the best of our knowledge, the only existing proposal that makes an effort in this direction is [15]. However,



Fig. 1. The conceptual components of db-nets

it employs workflow nets [1] for capturing the process control flow, without leveraging the advanced capabilities of CPNs. Taking inspiration from [15], we then propose db-nets, a new, balanced formal model for data-aware processes, rooted in CPNs and relational databases. We rigorously describe the abstractions offered by the model, and formalize its execution semantics. We finally invite the research community to build on this new model, discussing its potential along three subjects: modeling, verification, and simulation.

# 2 The DB-Net Model

In our formal model, called db-net, we combine the distinctive features of CPNs and relational databases into a coherent framework, sketched in Figure 1. The model is structured in three layers:

- *persistence layer*, capturing a full-fledged relational database with constraints, and used to store background data, and data that are persistent across cases.
- *control layer*, employing a variant of CPNs to capture the process control-flow, case data, and possibly the resources involved in the process execution.
- data logic layer, interconnecting in the persistence and the control layer.

Thanks to the data logic, the control layer is supported in querying the underlying persistent data and tunes its own behavior depending on the obtained answers. Furthermore, the data logic may be exploited by the control layer to update the persistent data depending on the current state, the data locally carried by tokens, and additional data obtained from the external world. We formalize the framework layer by layer, from the bottom to the top.

### 2.1 Persistence Layer

The persistence layer maintains the relevant data in the domain of interest. To this end, we rely on standard relational databases equipped with constraints, in the spirit of [9]. First-order (FO) constraints allow for the formalization of conventional database constraints, such as keys and functional dependencies, as well as semantic constraints reflecting the domain of interest. Differently from [9], though, we also consider data types, on the one hand resembling concrete logical schemas of relational databases (where table columns are typed), and on the other reconciling the persistence layer with the notion of "color" in CPNs. **Definition 1 (Data type, type domain).** A data type  $\mathcal{D}$  is a pair  $\langle \Delta_{\mathcal{D}}, \Gamma_{\mathcal{D}} \rangle$ , where  $\Delta_{\mathcal{D}}$  is a value domain, and  $\Gamma_{\mathcal{D}}$  is a finite set of predicate symbols. Each predicate symbol  $S \in \Gamma_{\mathcal{D}}$  comes with an arity  $n_S$  and an *n*-ary predicate  $S^{\mathcal{D}} \subseteq \Delta_{\mathcal{D}}^n$  that rigidly defines its semantics. A type domain is a finite set of data types.  $\Box$ 

In the following, we use  $\mathfrak{D}$  to denote a type domain of interest, assuming that types in  $\mathfrak{D}$  are pairwise disjoint, that is, their domains do not intersect, and their predicate symbols are syntactically distinguished. This guarantees that a predicate symbol S defined in some type of  $\mathfrak{D}$ , is defined only in that type, which can be then unambiguously denoted, with slight abuse of notation, by type(S). We also employ  $\Delta_{\mathfrak{D}} = \bigcup_{\mathcal{D} \in \mathfrak{D}} \Delta_{\mathcal{D}}$ . Examples of data types are:

- string :  $\langle \mathbb{S}, \{=_s\} \rangle$ , strings with the equality predicate;
- real :  $\langle \mathbb{R}, \{=_r, <_r\} \rangle$ , real numbers with the usual comparison operators;
- int: ⟨ℤ, {=<sub>int</sub>, <<sub>int</sub>, succ}⟩, integers with the usual comparison operators, as well as the successor predicate.

We employ Definition 1 to introduce a typed relational database.

**Definition 2 (Typed relation schema, arity).** A  $\mathfrak{D}$ -typed relation schema is a pair  $\langle R, \vec{\mathcal{D}} \rangle$ , where R is a relation name, and  $\vec{\mathcal{D}}$  is a tuple of elements from  $\mathfrak{D}$ , indicating the data types associated to each component of R. The number  $|\vec{\mathcal{D}}|$  is the arity of R.

**Definition 3 (Typed database schema).** A  $\mathfrak{D}$ -typed database schema  $\mathcal{R}$  is a finite set of  $\mathfrak{D}$ -typed relation schemas.

For compactness, we represent a typed relation schema  $\langle R, \langle \mathcal{D}_1, \ldots, \mathcal{D}_n \rangle \rangle$  using notation  $R(\mathcal{D}_1, \ldots, \mathcal{D}_n)$ . E.g.,  $Emp(\mathbf{string}, \mathbf{string}, \mathbf{real})$  models a ternary relation for employees, where the first component is a string denoting the employee id, the second a string for her name, and the third a real for her salary.

**Definition 4 (Typed database instance, active domain).** Given a  $\mathfrak{D}$ -typed database schema  $\mathcal{R}$ , a  $\mathfrak{D}$ -typed database instance  $\mathcal{I}$  over  $\mathcal{R}$  is a finite set of facts of the form  $R(\mathfrak{o}_1,\ldots,\mathfrak{o}_n)$ , such that (i)  $R(\mathcal{D}_1,\ldots,\mathcal{D}_n) \in \mathcal{R}$ , and (ii) for each  $i \in \{1,\ldots,n\}$ , we have  $\mathfrak{o}_i \in \Delta_{\mathcal{D}_i}$ . Given a type  $\mathcal{D} \in \mathfrak{D}$ , the  $\mathcal{D}$ -active domain of  $\mathcal{I}$ , written  $Adom_{\mathcal{D}}(\mathcal{I})$ , is the set of values in  $\Delta_{\mathcal{D}}$  such that  $\mathfrak{o} \in Adom_{\mathcal{D}}(\mathcal{I})$  if and only if  $\mathfrak{o} \in \Delta_{\mathcal{D}}$  and  $\mathfrak{o}$  occurs in  $\mathcal{I}$ .

We now turn to queries. As query language, we resort to standard first-order logic (FOL), interpreted under the active domain semantics [24]. This means that quantifiers are relativized to the active domain of the database instance of interest, guaranteeing that queries are domain-independent (actually, safe-range): their evaluation only depends on the values explicitly appearing in the database instance over which they are applied. Recall that this query language is equivalent to the well-known SQL standard [6]. Since the relational structures we consider are typed, the logic is typed as well.

Given a type domain  $\mathfrak{D}$ , we fix a countably infinite set  $\mathcal{V}_{\mathfrak{D}}$  of variables. Each variable is typed. To this end, we introduce a *variable typing function*  type:  $\mathcal{V}_{\mathfrak{D}} \to \mathfrak{D}$  mapping variables to their types. The typing function prescribes that x may be substituted only by values taken from  $\Delta_{type(x)}$ . For compactness, the variable type may be explicitly shown using a colon notation x:type(x).

**Definition 5 (FO(\mathfrak{D}) query).** A *(well-typed)* FO( $\mathfrak{D}$ ) *query* over a  $\mathfrak{D}$ -typed database schema  $\mathcal{R}$  is a formula of the form:

$$Q ::= S(\vec{y}) \mid R(\vec{z}) \mid \neg Q \mid Q_1 \land Q_2 \mid \exists x.Q, \text{ where}$$

- for  $\vec{y} = \langle y_1, \ldots, y_n \rangle$ , we have that S/n is a predicate defined in  $\Gamma_{\mathcal{D}}$  for some  $\mathcal{D} \in \mathfrak{D}$ , and for each  $i \in \{1, \ldots, n\}$ , we have that  $y_i$  is either a value  $\mathbf{o} \in \Delta_{\mathcal{D}}$ , or a variable  $x \in \mathcal{V}_{\mathfrak{D}}$  with  $\mathsf{type}(x) = \mathcal{D}$ ;
- for  $\vec{z} = \langle z_1, \ldots, z_m \rangle$ , we have that  $R(\mathcal{D}_1, \ldots, \mathcal{D}_m)$  is a relation defined in  $\mathcal{R}$ , and for each  $i \in \{1, \ldots, m\}$ , we have that  $z_i$  is either a value  $o \in \Delta_{\mathcal{D}_i}$ , or a variable  $x \in \mathcal{V}_{\mathfrak{D}}$  with  $type(x) = \mathcal{D}_i$ .

We use standard abbreviations  $Q_1 \lor Q_2 = \neg(\neg Q_1 \land \neg Q_2)$ , and  $\forall x.Q = \neg \exists x. \neg Q$ .

**Definition 6 (Free variable, boolean query).** A variable  $x \in \mathcal{V}_{\mathfrak{D}}$  is free in a FO( $\mathfrak{D}$ ) query Q, if x occurs in Q but is not in the scope of any quantifier. We use Free(Q) to denote the set of variables occurring free in Q. A boolean query is a query without free variables.

Given a query Q such that  $Free(Q) = \{x_1, \ldots, x_n\}$ , we employ notation  $\mathbb{Q}_{name}(x_1, \ldots, x_n)$ :- Q to emphasize the free variables of Q, and to fix a natural ordering over them. As usual, queries are used to extract answers from a database instance of interest.

**Definition 7 (Substitution).** Given a set  $X = \langle x_1, \ldots, x_n \rangle$  of typed variables, a *substitution* for X is a function  $\theta : X \to \Delta_{\mathfrak{D}}$  mapping variables from X into values, such that for every  $x \in X$ , we have  $\theta(x) \in \Delta_{type(x)}$ . A *substitution*  $\theta$  for  $a \ \mathsf{FO}(\mathfrak{D})$  query Q is a substitution for the free variables of Q.  $\Box$ 

As customary, we may view a substitution  $\theta$  for a query Q simply as a tuple of values, assuming the natural ordering over the free variables of Q. We denote by  $Q\theta$  the boolean query obtained from Q by replacing each free variable  $x \in Free(Q)$  with the corresponding value  $\theta(x)$ . In the following, we apply substitutions to any structure containing variables.

**Definition 8 (Query entailment with active domain semantics).** Given a  $\mathfrak{D}$ -typed database schema  $\mathcal{R}$ , a  $\mathfrak{D}$ -typed instance  $\mathcal{I}$  over  $\mathcal{R}$ , a  $FO(\mathfrak{D})$  query Qover  $\mathcal{R}$ , and a substitution  $\theta$  for Q, we inductively define the relation  $\mathcal{I}$  entails Q under  $\theta$  with active domain semantics, written  $\mathcal{I}, \theta \models Q$ , as:

$$\begin{array}{lll} \mathcal{I}, \theta \models R(y_1, \dots, y_n) \text{ if } & R(y_1, \dots, y_n) \theta \in \mathcal{I} \\ \mathcal{I}, \theta \models S(y_1, \dots, y_n) \text{ if } & S(y_1, \dots, y_n) \theta \in S^{\mathsf{type}(S)} \\ \mathcal{I}, \theta \models \neg Q & \text{ if } & \mathcal{I}, \theta \not\models Q \\ \mathcal{I}, \theta \models Q_1 \land Q_2 & \text{ if } & \mathcal{I}, \theta \models Q_1 \text{ and } \mathcal{I}, \theta \models Q_2 \\ \mathcal{I}, \theta \models \exists x. Q & \text{ if } & \text{there exists } \mathsf{o} \in Adom_{\mathsf{type}(x)}(\mathcal{I}) \text{ such that } \mathcal{I}, \theta[x/\mathsf{o}] \models Q \end{array}$$

where  $\theta[x/\mathfrak{o}]$  denotes the substitution obtained from  $\theta$  by assigning  $\mathfrak{o}$  to  $x^{1}$ 

<sup>&</sup>lt;sup>1</sup> If  $\theta(x)$  is defined, its value is replaced by  $\circ$ , otherwise  $\theta$  is extended so that  $\theta(x) = \circ$ .

**Definition 9 (Query answers).** Given a  $\mathfrak{D}$ -typed database schema  $\mathcal{R}$ , a  $\mathfrak{D}$ -typed instance  $\mathcal{I}$  over  $\mathcal{R}$ , and a  $\mathsf{FO}(\mathfrak{D})$  query  $Q(x_1, \ldots, x_n)$  over  $\mathcal{R}$ , the set of answers to Q in  $\mathcal{I}$ , written  $ans(Q, \mathcal{I})$ , is the set of substitutions  $\theta$  from the free variables of Q to the active domain of  $\mathcal{I}$ , such that Q holds in  $\mathcal{I}$  under  $\theta$ :

$$ans(Q,\mathcal{I}) = \left\{ \begin{array}{c} \theta : \{x_1, \dots, x_n\} \to \\ Adom_{\mathtt{type}(x_1)}(\mathcal{I}) \times \dots \times Adom_{\mathtt{type}(x_n)}(\mathcal{I}) \end{array} \middle| \mathcal{I}, \theta \models Q \right\} \quad \Box$$

When Q is boolean, we write  $ans(Q, \mathcal{I}) \equiv \text{true if } \langle \rangle \in ans(Q, \mathcal{I})$ , or  $ans(Q, \mathcal{I}) \equiv$ false if  $ans(Q, \mathcal{I}) = \emptyset$ . We are finally ready to define the persistence layer.

**Definition 10 (Persistence layer).** A  $\mathfrak{D}$ -typed *persistence layer* is a pair  $\langle \mathcal{R}, \mathcal{E} \rangle$  where: (*i*)  $\mathcal{R}$  is a  $\mathfrak{D}$ -typed database schema; (*ii*)  $\mathcal{E}$  is a finite set { $\Phi_1, ..., \Phi_k$ } of boolean FO( $\mathfrak{D}$ ) queries over  $\mathcal{R}$ , modeling *constraints over*  $\mathcal{R}$ .

The presence of constraints calls for a definition of which database instances are compliant by a given persistence layer, i.e., satisfy its constraints.

**Definition 11 (Compliant database instance).** Given a  $\mathfrak{D}$ -typed persistence layer  $\mathcal{P} = \langle \mathcal{R}, \mathcal{E} \rangle$  and a  $\mathfrak{D}$ -typed database instance  $\mathcal{I}$ , we say that  $\mathcal{I}$  complies with  $\mathcal{P}$  if: (i)  $\mathcal{I}$  is defined over  $\mathcal{R}$ ; (ii)  $\mathcal{I}$  satisfies all constraints in  $\mathcal{E}$ , that is,  $ans(\bigwedge_{\Phi \in \mathcal{E}} \Phi, \mathcal{I}) \equiv$ true.

**Example 1.** The persistence layer  $\mathcal{P} = \langle \mathcal{R}, \mathcal{E} \rangle$  is a fragment of an information system used by a company to handle the submission of tickets, and their management by employees.  $\mathcal{R}$  employs types **string** and **int** to define the following relation schemas:

- *Emp*(**string**) lists employee (names);
- *Ticket*(**int**, **string**) models ticket (identifiers) and their description;
- *Resp*(string, int) models which employees handle which tickets: *Resp*(e, 1) indicates that the employee named e is responsible for ticket number 1.
- Log(int, string, string) represents a log table storing information about all the tickets processed so far, also listing their responsible employees and their description.

The persistence layer is also equipped with a set of constraints over  $\mathcal{R}$ , expressing (primary) keys, foreign keys, functional dependencies, and multiplicity constraints. E.g., the ticket number provides the primary key for *Ticket*, the second component of *Resp* references the primary key of *Ticket*, and *each employee can handle at most one ticket at a time*. It is well-known that such constraints can be formalized in FO [6]. E.g., the latter constraint may be formalized as:  $\forall e, t_1, t_2.Resp(e, t_1) \land Resp(e, t_2) \rightarrow t_1 = t_2$ .

### 2.2 Data Logic Layer

The data logic layer provides a bidirectional "interface" to interact with a database instance complying with a persistence layer of interest. On the one hand, the data logic allows one to *extract* data from the database instance using queries. On the other hand, it allows one to *update* the database instance, adding and deleting possibly multiple facts at once, with a *transactional* semantics: if the new database instance obtained after the update is still compliant with the persistence layer, the update is *committed*, otherwise it is *rolled back*. This approach is in line with how database management systems operate in practice.

To query the database instance, we use  $FO(\mathfrak{D})$  queries as in Definition 5. To update the database instance, we instead resort to the literature on data-centric processes [33,11], where *actions* are typically used to apply CRUD (create-readupdate-delete) operations over a relational database. Specifically, we adopt a minimalistic approach, keeping the actions as simple as possible. The approach is inspired by the well-known STRIPS language for planning, which has been adopted also in for data-centric processes [5]. More sophisticated forms of actions, as those in [9], can be seamlessly introduced.

**Definition 12 (Action).** A (parameterized) action over a  $\mathfrak{D}$ -typed persistence layer  $\langle \mathcal{R}, \mathcal{E} \rangle$  is tuple  $\langle \mathbf{n}, \vec{p}, F^+, F^- \rangle$ , where: (i)  $\mathbf{n}$  is the action name; (ii)  $\vec{p}$  is a tuple of pairwise distinct typed variables from  $\mathcal{V}_{\mathfrak{D}}$ , denoting the action (formal) parameters. (iii)  $F^+$  and  $F^-$  respectively represent a finite set of  $\mathcal{R}$ -facts over  $\vec{p}$ , to be added to and deleted from the current database instance. Given a typed relation  $R(\mathcal{D}_1, \ldots, \mathcal{D}_n) \in \mathcal{R}$ , an R-fact over  $\vec{p}$  has the form  $R(y_1, \ldots, y_n)$ , such that for every  $i \in \{1, \ldots, n\}$ ,  $y_i$  is either a value  $\mathbf{o} \in \Delta_{\mathcal{D}_i}$ , or a variable  $x \in \vec{p}$ with  $\mathsf{type}(x) = \mathcal{D}_i$ . An  $\mathcal{R}$ -fact is an R-fact for some relation R from  $\mathcal{R}$ .

To access the different components of an action  $\alpha = \langle \mathbf{n}, \vec{p}, F^+, F^- \rangle$ , we use a dot notation:  $\alpha \cdot \mathbf{name} = \mathbf{n}, \alpha \cdot \mathbf{params} = \vec{p}, \alpha \cdot \mathbf{add} = F^+$ , and  $\alpha \cdot \mathbf{del} = F^-$ .

We now turn to the semantics of actions. Actions are executed by grounding their parameters to values. Given an action  $\alpha$  and a (parameter) substitution  $\theta$  for  $\alpha$ , we call *action instance*  $\alpha\theta$  the (ground) action resulting from  $\alpha$  by substituting its parameters with corresponding values, as specified by  $\theta$ .

**Definition 13 (Action instance application).** Let  $\mathcal{P} = \langle \mathcal{R}, \mathcal{E} \rangle$  be a  $\mathfrak{D}$ -typed persistence layer,  $\mathcal{I}$  be a  $\mathfrak{D}$ -typed database instance  $\mathcal{I}$  compliant with  $\mathfrak{D}$ ,  $\alpha$  be an action over  $\mathcal{P}$ , and  $\theta$  be a substitution for *action*-params. The *application* of  $\alpha\theta$  on  $\mathcal{I}$ , written  $\operatorname{apply}(\alpha\theta, \mathcal{I})$ , is a database instance over  $\mathcal{R}$  obtained as  $(\mathcal{I}\setminus F_{\alpha\theta}^{-})\cup F_{\alpha\theta}^{+}$ , where: (i)  $F_{\alpha\theta}^{-} = \bigcup_{R(\vec{y})\in\alpha \cdot \operatorname{del}} R(\vec{y})\theta$ ; (ii)  $F_{\alpha\theta}^{+} = \bigcup_{R(\vec{y})\in\alpha \cdot \operatorname{add}} R(\vec{y})\theta$ . We say that  $\alpha\theta$  can be successfully applied to  $\mathcal{I}$  if  $\operatorname{apply}(\alpha\theta, \mathcal{I})$  complies with  $\mathcal{P}$ .

The application of an action instance amounts to ground all the facts contained in the definition of the action as specified by the given substitution, then applying the update on the given database instance, giving priority to additions over deletions (this is a standard approach, which unambiguously handles the situation in which the same fact is asserted to be added and deleted).

The data logic simply exposes a set of queries and a set of actions that can be used by the control layer to obtain data from the persistence layer, and to induce updates on the persistence layer.

**Definition 14 (Data logic layer).** Given a  $\mathfrak{D}$ -typed persistence layer  $\mathcal{P}$ , a  $\mathfrak{D}$ -typed data logic layer over  $\mathcal{P}$  is a pair  $\langle \mathcal{Q}, \mathcal{A} \rangle$ , where: (i)  $\mathcal{Q}$  is a finite set of  $FO(\mathfrak{D})$  queries over  $\mathcal{P}$ ; (ii)  $\mathcal{A}$  is a finite set of actions over  $\mathcal{P}$ .

**Example 2.** We make the scenario of Example 1 operational, introducing a data logic layer  $\mathcal{L}$  over  $\mathcal{P}$ .  $\mathcal{L}$  exposes two queries to inspect the persistence layer:

•  $Q_e(e)$ :-  $Emp(e) \land \neg \exists t. Resp(e, t)$ , to extract *idle* employees;

•  $Q_t(t, d)$ :- Ticket(t, d), to extract tickets and their description.

In addition,  $\mathcal{L}$  provides three main functionalities to manipulate tickets in persistence layer: ticket registration, assignment/release, and logging. Such functionalities are realized through four actions (where, for simplicity, we blur the distinction between an action and its name). The registration of a new ticket is managed by an action REG that, given an integer t, and two strings e and d, (REG **params** =  $\langle t, e, d \rangle$ , simultaneously creates a ticket identified by t and described by d into the persistence layer, and assigns the employee identified by e to such ticket (thus making her *busy*):

 $REG \cdot del = \{ Emp(e, idle) \} REG \cdot add = \{ Ticket(t, d), Resp(e, t) \}$ 

Two specular actions ASSIGN and RELEASE are exposed to assign or release a ticket to/from an employee, making her busy or idle. Both actions take as input a string for the employee name and an integer for a ticket it (ASSIGN-params = RELEASE-params =  $\langle e, t \rangle$ ), and update *e* by removing or adding that *e* is responsible of *t*:

 $\text{RELEASE} \cdot \text{del} = \text{ASSIGN} \cdot \text{add} = \{ Resp(e, t) \} \qquad \text{RELEASE} \cdot \text{add} = \text{ASSIGN} \cdot \text{del} = \emptyset$ 

Finally, an action LOG with LOG params =  $\langle t, e, d \rangle$  is exposed to flush the information related to a ticket into a log table. The action erases all information about the ticket, and logs that it has been processed, also recalling its employee and description:

$$LOG \cdot del = \{ Ticket(t, d), Resp(e, t) \} \qquad LOG \cdot add = \{ Log(t, e, d) \}$$

### 2.3 Control Layer

The control layer employs a variant of CPNs to capture the process control flow, and how it interacts with an underlying persistence layer through the functionalities provided by the idata logic. The spirit is to conceptually ground CPNs by adopting a data-oriented approach. This is done by introducing dedicated constructs exploiting such functionalities, as well as simple, declarative patterns to capture the typical token consumption/creation mechanism of CPNs.

Before introducing the different constitutive elements of the control layer together with their graphical appearance, we fix some preliminary notions. We consider the standard notion of a *multiset*. Given a set A, the set of *multisets* over A, written  $A^{\oplus}$ , is the set of mappings of the form  $m : A \to \mathbb{N}$ . Given a multiset  $S \in A^{\oplus}$  and an element  $a \in A$ ,  $S(a) \in \mathbb{N}$  denotes the number of times a appears in S. Given  $a \in A$  and  $n \in \mathbb{N}$ , we write  $a^n \in S$  if S(a) = n. We also consider the usual operations on multisets. Given  $S_1, S_2 \in A^{\oplus}$ : (i)  $S_1 \subseteq S_2$  (resp.,  $S_1 \subset S_2$ ) if  $S_1(a) \leq S_2(a)$  (resp.,  $S_1(a) < S_2(a)$ ) for each  $a \in A$ ; (ii)  $S_1 + S_2 = \{a^n \mid a \in A \text{ and } n = S_1(a) + S_2(a)\}$ ; (iii) if  $S_1 \subseteq S_2, S_2 - S_1 = \{a^n \mid a \in A \text{ and } n = S_2(a) - S_1(a)\}$ ; (iv) given a number  $k \in \mathbb{N}, k \cdot S_1 = \{a^{kn} \mid a^n \in S_1\}$ .<sup>2</sup>

**Places.** The control layer contains a finite set P of places, which in turn are classified in two groups. On the one hand, so-called *control places* play the role of standard places in classical Petri nets: they represent conditions/states of a

<sup>&</sup>lt;sup>2</sup> Hence, given a multiset S, we have  $0 \cdot S = \emptyset$ .

dynamic system. On the other hand, so-called view places are used as an interface to the underlying persistence layer, so as to make the persistent data available to the control layer. We then have  $P = P_c \uplus P_v$ , where  $P_c$  and  $P_v$  respectively denote the set of control and view places. Graphically, we depict control places using the standard notation:  $\bigcirc$ . We instead decorate view places as:  $\bigcirc$ .

In the spirit of CPNs, the control layer assigns to each place a color, which in turn combines one or more data types from a type domain  $\mathfrak{D}$ . Formally, a  $\mathfrak{D}$ -color is a cartesian product  $\mathcal{D}_1 \times \ldots \times \mathcal{D}_m$ , where for each  $i \in \{1, \ldots, m\}$ , we have  $\mathcal{D}_i \in \mathfrak{D}$ . We denote by  $\Sigma$  the set of all possible  $\mathfrak{D}$ -colors.

# **Definition 15 (Color assignment).** A $\mathfrak{D}$ -color assignment over places P is a function color : $P \to \Sigma$ mapping each place $p \in P$ to a corresponding $\mathfrak{D}$ -color.

As for control places, it is well-known that the coloring mechanism can be exploited to realize a plethora of conceptual abstractions on top of the control flow. We mention here the two most important abstractions in our setting: (i) cases and their data, and (ii) resource. A case represents a specific process instance, and its case data [31] are local data whose scope is the case itself, and that are used to store important information for the progression of the case. Such data may be either extracted from the underlying persistence layer, or obtained by interacting with the external environment (e.g., human users, external services, or data generators). Resources represent actors able to handle the execution of tasks. They are also typically associated to data attributes (e.g., id, role, group). Tasks typically consume (certain kinds of) resources when executed, and this implicitly affect the degree of concurrency in the progression of cases, as well as the possibility of spawning new cases.

The fact that control places are colored implies that whenever a token is assigned to a control place, it must carry a data tuple whose types match component-wise the place color. It is worth noting that a colored place may be interchangeably considered as a specific state/condition within the control layer, or as a special relation schema used to enrich the persistence layer with control-related information. Similarly, a token distributed over a place may be interchangeably seen as a thread of control located in that state, or as a tuple assigned to the relation schema represented by that place.

As discussed above, control places host tokens carrying local data. Obviously, the control layer also requires to query persistent data, using them to decide how to route tokens when it comes to business decisions, or to assign them to case data. We want to support both possibilities, but clearly separating the data retrieved from the persistence layer, from those carried by tokens. This is why we distinguish view places from control places. Each view place exposes to the control layer a portion of the data stored in the persistence layer. Formally, this is done by equipping the view place with a query defined in the data logic layer.

**Definition 16 (Query assignment).** Given a data logic layer  $\mathcal{L} = \langle \mathcal{Q}, \mathcal{A} \rangle$ , a query assignment from view places  $P_v$  to queries  $\mathcal{Q}$  is a function query :  $P_v \to \mathcal{Q}$  mapping each view place  $p \in P_v$  with  $color(p) = \mathcal{D}_1 \times \ldots \times \mathcal{D}_n$  to a query

 $Q(x_1, \ldots, x_n)$  from  $\mathcal{Q}$ , such that the color of p component-wise matches with the types of the free variables in Q: for each  $i \in \{1, \ldots, n\}$ , we have  $\mathcal{D}_i = type(x_i).\Box$ 

A view place may be seen as a normal place, whose color is implicitly obtained by the types of the free variables of the query, considered with their natural ordering. However, tokens are not arbitrarily attached to it: at a given time, the tokens it contains represent the answers to the query it is associated to. All such tokens are only "virtually" present in the control layer, and in fact they cannot be consumed within the control layer itself, but only accessed in a read-only way. Notice, however, that the content of the view place is not immutable: it changes whenever the data it fetches from the persistence layer are updated.

**Transitions.** As customary, in our model transitions represent atomic units of work within the control layer, thus providing the fundamental building block to describe the dynamics of a process. As usual, they are depicted using a square notation:  $\Box$ . In our setting, they simultaneously account for three different aspects: the token consumption/production mechanism of CPNs, the injection of possibly fresh data from the external environment a là  $\nu$ -Petri nets [30], and the impact on the underlying persistence layer.

We start with the consumption of tokens. This is modeled through input arcs connecting places to transitions, together with inscriptions that declaratively match tokens and their data. To this end, we build on the approach adopted in variants of data nets [30,23,5]: an inscription is just a multiset of tuples over a given set of typed variables. Each tuple matches with a token present in the input place, and the variables therein are bound, component-wise, to the data carried by such a token. In addition, if the input place is a control place, the token is consumed upon firing, whereas if the place is a view place, it is only inspected. Graphically, we adopt the following conventions. An input arc from a control place is depicted as usual:  $\bigcirc \frown \square$ . An input arc from a view place is instead depicted using the read-arc notation:

The overall consumption/inspection of tokens and the data they carry along all arcs incoming into a transition constitutes a *firing mode* for that transition. In the context of a transition definition, we call a tuple of typed variables (as well as, possibly, values) *inscription*. We denote the set of all possible inscriptions over set  $\mathcal{Y}$  as  $\Omega_{\mathcal{Y}}$ . In addition, we denote the set of variables appearing inside an inscription  $\omega \in \Omega_{\mathcal{Y}}$  as  $Vars(\omega)$ , and we extend such notation to sets and multisets of inscriptions.

**Definition 17 (Input flow).** An *input flow* from places P to transitions T is a function  $F_{in}: P \times T \to \Omega^{\oplus}_{\mathcal{V}_{\mathfrak{D}}}$  assigning multisets of inscriptions (over variables  $\mathcal{V}_{\mathfrak{D}}$ ) to input arcs, such that all such inscriptions are compatible with their input places. An inscription  $\langle x_1, \ldots, x_m \rangle$  is *compatible* with a place p if  $color(p) = \mathcal{D}_1 \times \ldots \times \mathcal{D}_m$ , such that for every  $i \in \{1, \ldots, m\}$ , we have  $type(x_i) = \mathcal{D}_i$ .  $\Box$ 

Graphically, we do not depict input arcs whose inscription is  $\emptyset$ . We define the *input variables* of t, written InVars(t) as the set of all variables occurring on input arc inscriptions for t:

 $InVars(t) = \{x \in \mathcal{V}_{\mathfrak{D}} \mid \text{there exists } p \in P \text{ such that } x \in Vars(F_{in}(\langle p, t \rangle))\}.$ 

The set InVars(t) gives an indication about which input data elements are accessed when a transition fires. The multiple usage of the same variable in an inscription, or in inscriptions attached to different arcs incident to a transition, captures the requirement of *matching* the same data object in different tokens, allowing the transition to fire only if the accessed tokens carry the *same* data value. This mirrors the notion of join used when querying relational data. In general, though, the modeler may require to specify additional constraints over such input data to allow firing the transition. To this end, we introduce guards.

**Definition 18 (Guard).** A  $\mathfrak{D}$ -typed guard is a formula of the form:

$$\varphi ::= \operatorname{true} | S(\vec{y}) | \neg \varphi | \varphi_1 \land \varphi_2$$

where, for  $\vec{y} = \langle y_1, \ldots, y_n \rangle \subseteq \mathcal{V}_{\mathfrak{D}}$ , we have that S/n is a predicate defined in  $\Gamma_{\mathcal{D}}$  for some  $\mathcal{D} \in \mathfrak{D}$ , and for each  $i \in \{1, \ldots, n\}$ , we have that  $y_i$  is either a value  $o \in \Delta_{\mathcal{D}}$ , or a variable  $x_i \in \mathcal{V}_{\mathfrak{D}}$  with  $type(x_i) = \mathcal{D}$ .

We denote by  $\mathbb{F}_{\mathfrak{D}}$  the set of all possible  $\mathfrak{D}$ -typed guards. Additionally, with a slight abuse of notation, given guard  $\varphi$  we denote by  $Vars(\varphi)$  the set of variables occurring in  $\varphi$ . Guards may be seen as the quantifier- and relation-free fragment of  $FO(\mathfrak{D})$  queries (cf. Definition 5). Consequently, their semantics is inherited from Definition 8 (considering the empty database instance). Guards are attached to transitions, and defined over their input variables, thus being an additional filter on the data that can be matched to the input inscriptions.

**Definition 19 (Transition guard assignment).** A  $\mathfrak{D}$ -typed transition guard assignment over transitions T is a function guard :  $T \to \mathbb{F}_{\mathfrak{D}}$  assigning to each transition  $t \in T$  a  $\mathfrak{D}$ -typed guard  $\varphi$ , such that  $Vars(\varphi) \subseteq InVars(t)$ .

We now concentrate on the effect of firing a transition, which may simultaneously impact the control layer and the underlying persistence layer. Such an effect is tuned by the input variables attached to the transition, as well as additional data obtained from the external environment. Injection of external data is crucial for two reasons [11,26,5]. First, during the execution of a case, input data may be dynamically acquired from human users or external services, and used later on; this is, e.g., what happens when a user form needs to be filled before continuing with the case execution, then deciding how to route the case depending on the inserted data. Second, fresh identifiers may be injected into the system, e.g., to explicitly distinguish tokens via certain data attributes, or to insert a new tuple in the underlying database instance (which typically requires to create a distinctive primary key for that tuple). We call these two types of external inputs arbitrary external inputs and fresh external inputs. To account for arbitrary external inputs in the context of a transition, we just employ "normal" variables distinct from those used in the input inscriptions. To account for fresh external inputs, we employ the well-known mechanism adopted in  $\nu$ -Petri nets [30,27]. In particular, we introduce a countably infinite set  $\Upsilon_{\mathfrak{D}}$  of  $\mathfrak{D}$ -typed fresh

variables. To guarantee an unlimited provisioning of fresh values, we impose that for every variable  $\nu \in \Upsilon_{\mathfrak{D}}$ , we have that  $\Delta_{type(\nu)}$  is countably infinite.

From now on, we fix a countably infinite set of  $\mathfrak{D}$ -typed variable  $\mathcal{X}_{\mathfrak{D}}$ , obtained as the disjoint union of "normal" variables  $\mathcal{V}_{\mathfrak{D}}$  and fresh variables  $\mathcal{T}_{\mathfrak{D}}$ . In formulae,  $\mathcal{X}_{\mathfrak{D}} = \mathcal{V}_{\mathfrak{D}} \uplus \mathcal{T}_{\mathfrak{D}}$ . Let us first focus on the impact of transition firing on the underlying persistence layer. This is, again, mediated by the data logic, exploiting in particular the actions it exposes. Specifically, a transition can bind to an action, using variables from  $\mathcal{X}_{\mathfrak{D}}$  as "actual" parameters. In this light, data passing from the control to the persistence layer is captured by re-using the same variable inside an input inscription and an action binding for the same transition. When the transition fires, actual parameters are substituted with concrete data values, instanating the action and allowing for its further invocation.

**Definition 20 (Action assignment).** Given a data logic layer  $\mathcal{L} = \langle \mathcal{Q}, \mathcal{A} \rangle$ , an *action assignment* from transitions T to actions  $\mathcal{A}$  is a partial function  $\operatorname{act} : T \to \mathcal{A} \times \Omega_{\mathcal{X}_{\mathcal{D}} \cup \mathcal{\Delta}_{\mathcal{D}}}$ , where  $\operatorname{act}(t)$  maps t to an action  $\alpha \in \mathcal{A}$  together with a (binding) inscription compatible with  $\alpha$ . An inscription  $\langle y_1, \ldots, y_m \rangle$  is compatible with  $\alpha$  if  $\alpha$ -params =  $\langle z_1, \ldots, z_m \rangle$  and, for each  $i \in \{1, \ldots, m\}$ , we have  $\operatorname{type}(y_i) = \operatorname{type}(z_i)$  if  $y_i$  is a variable from  $\mathcal{X}_{\mathfrak{D}}$ , or  $y_i \in \mathcal{\Delta}_{\operatorname{type}(z_i)}$  if  $y_i$  is a value from  $\mathcal{\Delta}_{\mathfrak{D}}$ .

The action assignment provides a distinctive feature of our model, namely the ability of the control layer to invoke an action applied to the underlying persistence layer. This, however, does not in general guarantee that the action invocation will actually turn into an update over the persistence layer. Recall in fact that an action instance is applied transactionally: if it produces a new database instance that is compliant with the persistence layer, the action instance *succeeds* and the update is committed; if, instead, some constraints is violated, the action instance *fails* and the update does not take place.

Lastly, we consider the effect of transitions on the control layer itself, defining which tokens have to produced, together with the data they will carry, and to which places such tokens have to be assigned. This is done by mirroring the definition of input flow (cf. Definition 17), with two distinctions. First, output arcs connect transitions to control places only, as view places cannot be explicitly modified within the control layer. Second, the inscriptions attached to output arcs may mention not only input variables, but also: (i) values, allowing for constructing tokens that carry explicitly specified data; (ii) fresh variables, allowing for constructing tokens that carry data not already present in the net, nor in the underlying database instance.

**Definition 21 (Output flow).** An *output flow* from transitions T to control places  $P_c$  is a function  $F_{out} : T \times P_c \to \Omega^{\oplus}_{\mathcal{X}_{\mathfrak{D}} \cup \mathcal{\Delta}_{\mathfrak{D}}}$  assigning multisets of inscriptions to output arcs, such that all such inscriptions are compatible with their output places (as defined in Definition 17).

We do not depict output arcs graphically when their inscription is  $\emptyset$ . We define the *output variables* of t, written *OutVars*(t), as the set of variables occurring in

the action assignment for t (if any), and in its output arc inscriptions:

 $\begin{aligned} OutVars(t) &= \{ x \in \mathcal{X}_{\mathfrak{D}} \mid \mathtt{act}(t) \text{ is defined as } \langle \alpha, \omega \rangle, \text{ and } x \in Vars(\omega) \} \\ & \cup \{ x \in \mathcal{X}_{\mathfrak{D}} \mid \text{there exists } p \in P \text{ such that } x \in Vars(F_{out}(\langle t, p \rangle)) \}. \end{aligned}$ 

With this notion at hand, we can obtain the external variables of transition t as  $OutVars(t) \setminus InVars(t)$ . Each such variable x is not bound by any input inscription, and can consequently be assigned arbitrarily (if  $x \in \mathcal{V}_{\mathfrak{D}}$ ), or to a fresh value (if  $x \in \Upsilon_{\mathfrak{D}}$ ). Among such variables, we explicitly refer to the fresh variables attached to t, using notation FreshVars(t). Mathematically,  $FreshVars(t) = OutVars(t) \cap \Upsilon_{\mathfrak{D}}$ . As discussed before, firing a transition may incur in the instantiation and invocation of an action from the data logic layer, and the so-obtained action instance may or not result in an actual update. To raise awareness of the control layer about these two radically different outcomes, we introduce two separate output flows: a normal output flow, capturing the actual effect of a transition on the control flow when its attached action succeeds, and a rollback flow, capturing the actual effect of a transition on the control flow when its attached action fails. With this distinction, the control layer can fine-tune its own behavior in accordance with the transactional semantics of the persistence layer, e.g., taking a standard or a compensation route depending on the outcome of the action. To graphically distinguish normal output arcs from rollback output arcs, we proceed as follows. We depict the former as usual:  $\Box \longrightarrow O$ . Instead, we decorate the latter with an "x":  $\square \longrightarrow \bigcirc$ .

We are finally in the position of defining the control layer.

**Definition 22 (Control layer).** A  $\mathfrak{D}$ -typed *control layer* over a data logic layer  $\mathcal{L} = \langle \mathcal{Q}, \mathcal{A} \rangle$  is a tuple  $\langle P, T, F_{in}, F_{out}, F_{rb}, \text{color}, \text{query}, \text{guard}, \text{act} \rangle$ , where:

- $P = P_c \uplus P_v$  is a finite set of control places constituted by control places  $P_c$ and view places  $P_v$ ;
- T is a finite set of transitions, such that  $T \cap P = \emptyset$ ;
- $F_{in}$  is an input flow from P to T (cf. Definition 17);
- $F_{out}$  and  $F_{rb}$  are two output flows from T to  $P_c$  (cf. Definition 21), respectively called *normal output flow* output flows flow;
- color is a color assignment over P (cf. Definition 15);
- query is a query assignment from  $P_v$  to  $\mathcal{Q}$  (cf. Definition 16);
- guard is a transition guard assignment over T (cf. Definition 19);
- act is an action assignment from T to  $\mathcal{A}$  (cf. Definition 20).

#### 2.4 DB-nets

We now put the three layers together, providing a formal definition for db-nets.

**Definition 23 (Db-net).** A *db-net* is a tuple  $\langle \mathfrak{D}, \mathcal{P}, \mathcal{L}, \mathcal{N} \rangle$ , where:

- $\mathfrak{D}$  is a type domain (cf. Definition 1);
- $\mathcal{P}$  is a  $\mathfrak{D}$ -typed persistence layer (cf. Definition 10);
- $\mathcal{L}$  is a  $\mathfrak{D}$ -typed data logic layer over  $\mathcal{P}$  (cf. Definition 14);



**Fig. 2.** The control layer of a db-net for ticket management. In CreateTicket,  $\nu_t$  is a fresh input variable, and *descr* is an arbitrary input variable.

•  $\mathcal{N}$  is a  $\mathfrak{D}$ -typed control layer over  $\mathcal{L}$  (cf. Definition 22).

We close the section by equipping our running example with a control layer.

**Example 3.** Figure 2 shows the control layer of a db-net  $\mathcal{B}$ , using the persistence layer  $\mathcal{P}$  defined in Example 1 and the data logic layer  $\mathcal{L}$  defined in Example 2.2. The control layer realizes a simple ticket processing workflow, where tickets are created, manipulated, and finally resolved. In spite of its simplicity,  $\mathcal{B}$  already shows many distinctive features of our model. We intuitively describe the control layer moving from left to right and from top to bottom. Each case of this process is constituted by a ticket and its responsible employee. A ticket is created by the CreateTicket transition, which requires the presence of an idle employee to be fired. Since this condition needs to inspect the persistence layer so as to retrieve idle employees, we model it through a view place associated to query  $Q_e$  from  $\mathcal{L}$ . Notice that if no employee is currently idle, then CreateTicket is not enabled. Upon firing CreateTicket for a given idle employee, a fresh ticket identifier is generated using fresh variable  $\nu_t$ , and a ticket description is obtained through the "external" input variable descr. All such data are bound to action REGISTER, which is applied when the transition fires. Among the effects of REGISTER, there is one asserting that the selected employee becomes responsible for the newly created ticket. This indirectly implies that such an employee is not present anymore in the view place for idle employees. The ticket id, together with its responsible employee, represent the case and its data. The two control places Active Tickets and Stalled Tickets have color  $\mathbf{int} \times \mathbf{string}$ , and model two distinct states in which tickets may be. Such states are important only within the evolution of cases, and are therefore not propagated to the underlying persistence layer. An active ticket may be "stalled" if the employee is currently unable to resolve it. Executing the stall transition has a twofold effect. Within the control layer, the ticket is moved from active to stalled. Within the persistence layer, its responsible employee is released. Interestingly, the relation of responsibility is now only recalled within the control layer. A stalled ticket may be revived, by inserting such a relation back into the persistence layer. This is captured by the Awake transition, which mirrors the effect of the Stall transition. However, there is a particularly interesting aspect here. When a ticket  $t_1$  is stalled, its responsible employee e is released and becomes idle. She may be then selected as responsible of a newly created ticket  $t_2$ . Due to the constraints present in  $\mathcal{P}$ , the indirect effect of this situation is that  $t_1$  cannot be awaken unless  $t_2$  is either stalled or resolved. In fact, awakening  $t_1$  in a situation where  $t_2$  is active would violate the requirement that e is

responsible of at most one ticket. For this reason, we enrich the Awake transition with a rollback output arc, which brings back the ticket to the stalled state if it is awaken in the "wrong" moment. For example, if  $t_1$  is awaken while  $t_2$  is active, the application of ASSIGN applied to  $\langle t_1, e \rangle$  will fail, consequently bringing  $t_1$  back to stalled. Finally, an active ticket may be resolved. This has a twofold effect. On the one hand, the token carrying the ticket and its responsible employee is removed from the net. On the other hand, the case information is logged into the persistence layer. However, logging also requires to retrieve the description of the ticket. To this end, we employ a second view place accessing tickets and their description by exploiting  $Q_t$  from  $\mathcal{L}$ . By using the same variable *tid* in the two input inscriptions of the Resolve transition, we realize a join, thus inspecting the view place and extracting the description of *tid*.

## **3** Execution Semantics

The execution semantics of a db-net simultaneously accounts for the progression of a database instance compliant with the persistence layer of the net, and for the evolution of a marking over the control layer of the net. Such two information sources affect each other via the data logic layer: the database instance exposes its own data through view places, influencing the current marking and the enabled transitions; the marking over the control layer determines which transitions may be fired, in turn triggering updates the database instance. We start by formalizing the notion of marking over the control layer of the db-net. A marking distributes tokens over the places of the net, so that each token carries data that are compatible with the color of the place in which that token resides. In this light, tokens are nothing else than tuples of values over the place colors. In addition, the marking of a view place must correspond to the answers obtained by issuing its associated query over the underlying database instance.

**Definition 24 (Marking).** A marking of a  $\mathfrak{D}$ -typed control layer  $\mathcal{N} = \langle P, T, F_{in}, F_{out}, F_{rb}, \text{color}, \text{query}, \text{guard}, \text{act} \rangle$  is a function  $m: P \to \Omega_{\mathfrak{D}}^{\oplus}$  mapping each place  $p \in P$  to a corresponding multiset of *p*-compatible tuples using data values from  $\mathfrak{D}$ . A tuple  $\langle \mathsf{o}_1, \ldots, \mathsf{o}_n \rangle$  is *p*-compatible if  $\mathsf{color}(p)$  is of the form  $\langle \mathcal{D}_1, \ldots, \mathcal{D}_n \rangle$ , and for every  $i \in \{1, \ldots, n\}$ , we have  $\mathsf{o}_i \in \Delta_{\mathcal{D}_i}$ . Given a database instance  $\mathcal{I}$ , we say that *m* is aligned to  $\mathcal{I}$  via query if the tuples it assigns to view places exactly correspond to the answers of their corresponding queries over  $\mathcal{I}$ : for every view place  $v \in P$  and every *v*-compatible tuple  $\vec{o}$ , we have that  $\vec{o} \in m(v)$  if and only if  $\vec{o} \in ans(query(v), \mathcal{I})$ .

We mirror the notion of active domain as provided in Definition 4 to the case of markings. Given a type  $\mathcal{D} \in \mathfrak{D}$ , the  $\mathcal{D}$ -active domain of a marking m, written  $Adom_{\mathcal{D}}(m)$ , is the set of values in  $\Delta_{\mathcal{D}}$  such that  $\mathfrak{o} \in Adom_{\mathcal{D}}(m)$  if and only if there exists p such that  $\mathfrak{o}$  occurs in m(p). From the practical point of view, one may consider the marking of control places to be initially defined by the modeler, and then evolved by the control layer, while the marking of view places computed on-the-fly from the underlying database instance when needed.

In db-nets, then, both the persistence layer and the control layer are stateful: during the execution, the persistence layer is associated to a database instance, while the control layer to a marking aligned with that database instance.

**Definition 25 (Snapshot).** Given a db-net  $\mathcal{B} = \langle \mathfrak{D}, \mathcal{P}, \mathcal{L}, \mathcal{N} \rangle$  with control layer  $\mathcal{N} = \langle P, T, F_{in}, F_{out}, F_{rb}, \text{color}, \text{query}, \text{guard}, \text{act} \rangle$ , a *snapshot* of  $\mathcal{B}$  (also called  $\mathcal{B}$ -snapshot) is a pair  $\langle \mathcal{I}, m \rangle$ , where  $\mathcal{I}$  is a database instance compliant with  $\mathcal{P}$ , and m is a marking of  $\mathcal{N}$  aligned to  $\mathcal{I}$  via query.

As customary for CPNs, the firing of a transition t in a snapshot is defined w.r.t. a so-called *binding* for t, that is, a substitution  $\sigma : Vars(t) \to \Delta_{\mathfrak{D}}$ , where  $Vars(t) = InVars(t) \cup OutVars(t)$ . However, to properly enable the firing of t, the binding  $\sigma$  must guarantee a number of properties:

- 1. agreement with the distribution of tokens over the places, in accordance with the inscriptions on the corresponding input arcs;
- 2. satisfaction of the guard attached to t;
- 3. proper treatment of fresh variables, guaranteeing that they are substituted with values that are pairwise distinct, and also distinct from all the values present in the current marking, as well as in the current database instance.

To formalize these conditions, we need the preliminary notion of inscription binding. Given an inscription (i.e., multiset of tuples of variables)  $\omega \in \Omega^{\oplus}_{\mathcal{X}_{\mathfrak{D}} \cup \mathcal{\Delta}_{\mathfrak{D}}}$ , and a substitution  $\theta$  defined over a set X of variables containing all variables occuring in  $\omega$ , the *inscription binding* of  $\omega$  under  $\theta$  is a multiset  $\theta^{\oplus}(\omega)$  from  $\Omega^{\oplus}_{\mathfrak{D}}$  defined as follows:  $\langle \mathsf{o}_1, \ldots, \mathsf{o}_n \rangle^m \in \theta^{\oplus}(\omega)$  if and only if  $\langle y_1, \ldots, y_n \rangle^m \in \omega$ , such that for every  $i \in \{1, \ldots, n\}$ , we have  $\mathsf{o}_i = y_i$  if  $y_i \in \mathcal{\Delta}_{\mathfrak{D}}$ , or  $\mathsf{o}_i = \theta(y_i)$  if  $y_i \in \mathcal{X}_{\mathfrak{D}}$ . For example, given  $\omega = \{\langle x, y \rangle^2, \langle x, 1 \rangle\}$  and  $\theta = \{x \mapsto 1, y \mapsto 2\}$ , we have  $\theta^{\oplus}(\omega) = \{\langle 1, 2 \rangle^2, \langle 1, 1 \rangle\}$ .

**Definition 26 (Transition enablement).** Let  $\mathcal{B}$  be a db-net with control layer  $\langle P, T, F_{in}, F_{out}, F_{rb}, \text{color}, \text{query}, \text{guard}, \text{act} \rangle$ . A transition  $t \in T$  is *enabled* in a  $\mathcal{B}$ -snapshot  $\langle \mathcal{I}, m \rangle$ , written  $\langle \mathcal{I}, m \rangle[t, \sigma \rangle$ , if:

- 1. for every place  $p \in P$ , m(p) provides enough tokens matching those required by inscription  $\omega = F_{in}(\langle p, t \rangle)$  once  $\omega$  is bound by  $\sigma$ , i.e.,  $\sigma^{\oplus}(\omega) \subseteq m(p)$ ;
- 2.  $guard(t)\sigma$  is true;
- 3.  $\sigma$  is injective over FreshVars(t), thus guaranteeing that fresh variables are assigned to pairwise distinct values by  $\sigma$ , and for every fresh variable  $\nu \in$  $FreshVars(t), \sigma(\nu) \notin (Adom_{type(\nu)}(\mathcal{I}) \cup Adom_{type(\nu)}(m))$ .  $\Box$

**Definition 27 (Induced action instance).** Let  $\mathcal{N} = \langle P, T, F_{in}, F_{out}, F_{rb},$ color, query, guard, act be a  $\mathfrak{D}$ -typed control layer, and let  $t \in T$  be a transition of  $\mathcal{N}$  such that  $\operatorname{act}(t) = \langle \alpha, \omega \rangle$ , with  $\alpha$ -params =  $\langle x_1, \ldots, x_n \rangle$  and  $\omega = \langle y_1, \ldots, y_n \rangle$ . The action instance induced by transition  $t \in T$  under binding  $\sigma$ , written  $\operatorname{act}_{\sigma}(t)$ , is the action instance  $\alpha \sigma'$ , where  $\sigma' : \alpha$ -params  $\rightarrow \mathcal{A}_{\mathfrak{D}}$  is a substitution for the formal parameters of  $\alpha$ , defined as: for every  $i \in \{1, \ldots, n\}$ , if  $y_i \in \mathcal{A}_{\mathfrak{D}}$ , then  $\sigma'(x_i) = y_i$ ; if instead  $y_i \in \mathcal{X}_{\mathfrak{D}}$ , then  $\sigma'(x_i) = \theta(y_i)$ . The firing of an enabled transition under some mode has then a threefold effect. First, all tokens present in control places that are used to match the input inscriptions are consumed. Second, the action instance induced by the firing is applied on the current database instance. If such an action instance can be successfully applied, the database instance is updated accordingly; if not, the database instance is kept unaltered (thus realizing a rollback). Third, tokens constructed using the inscriptions on output arcs are produced, and inserted into their target places, considering either normal output arcs or rollback arcs depending on whether the induced action instance is successfully applied or not.

**Definition 28 (Transition firing).** Let  $\mathcal{B} = \langle \mathfrak{D}, \mathcal{P}, \mathcal{L}, \mathcal{N} \rangle$  be a db-net with  $\mathcal{N} = \langle P, T, F_{in}, F_{out}, F_{rb}, \text{color}, \text{query}, \text{guard}, \text{act} \rangle$ , and  $s_1 = \langle \mathcal{I}_1, m_1 \rangle$ ,  $s_2 = \langle \mathcal{I}_2, m_2 \rangle$  be two  $\mathcal{B}$ -snapshots. Let  $t \in T$  be a transition of  $\mathcal{N}$ , and  $\sigma$  be a binding for t, such that  $s_1[t, \sigma]$ . We say that t fires in  $s_1$  with binding  $\sigma$  producing  $s_2$ , written  $s_1[t, \sigma]s_2$ , if the following conditions hold: given  $\mathcal{I}_3 = \text{apply}(\text{act}_{\sigma}(t), \mathcal{I}_1)$ ,

- if  $\mathcal{I}_3$  is compliant with  $\mathcal{P}$ , then  $\mathcal{I}_2 = \mathcal{I}_3$ , otherwise  $\mathcal{I}_2 = \mathcal{I}_1$ ;
- For every control place  $p \in P$ , given  $\omega_{in} = F_{in}(\langle p, t \rangle)$ ,  $\omega_{out} = F_{out}(\langle t, p \rangle)$ , and  $\omega_{rb} = F_{rb}(\langle t, p \rangle)$ , we have

$$m_2(p) = (m_1(p) - \sigma^{\oplus}(\omega_{in})) + k_{out} \cdot \sigma^{\oplus}(\omega_{out}) + (1 - k_{out}) \cdot \sigma^{\oplus}(\omega_{rb}),$$

where  $k_{out} = 1$  if  $\mathcal{I}_3$  is compliant with  $\mathcal{P}$ , and  $k_{out} = 0$  otherwise.

The execution semantics of a db-net is captured by a possibly *infinite-state* labeled transition system (LTS) that accounts for all possible executions of the control layer starting from an initial snapshot. States of this transition systems are db-net snapshots, and transitions model the effect of firing db-net transitions under given bindings. Formally, the execution semantics of a db-net  $\mathcal{B} = \langle \mathfrak{D}, \mathcal{P}, \mathcal{L}, \mathcal{N} \rangle$  with  $\mathcal{N} = \langle P, T, F_{in}, F_{out}, F_{rb}, \text{color}, \text{query}, \text{guard}, \text{act} \rangle$  is given in terms of an LTS  $\Gamma_{\mathcal{B}} = \langle S, s_0, \rightarrow \rangle$ , where:

- S is a possibly infinite set of  $\mathcal{B}$ -snapshots;
- $s_0$  is the *initial*  $\mathcal{B}$ -snapshot;
- $\rightarrow \subseteq S \times T \times S$  is a transition relation over states, labeled by transitions T;
- S and  $\rightarrow$  are defined by simultaneous induction as the smallest sets such that: (i)  $s_0 \in S$ ; (ii) given a  $\mathcal{B}$ -snapshot  $s \in S$ , for every transition  $t \in T$ , binding  $\sigma$ , and  $\mathcal{B}$ -snapshot s', if  $s[t, \sigma\rangle s'$  then  $s' \in S$  and  $s \stackrel{t}{\rightarrow} s'$ .

### 4 Discussion and Conclusion

We believe that db-nets have the potential of stimulating discussion, and possibly spawning new lines of investigation, for researchers interested in the foundations and applications of data-aware processes. We consider in particular the impact on modeling, verification, and simulation.

**Modeling.** From the modeling point of view, db-nets incorporate all typical abstractions needed in data-aware business processes. In this light, our model covers all the distinctive features of various Petri net classes enriched with data, as well as those of data-centric processes. More formally, in terms of expressiveness, we observe/conjecture the following correspondences. First of all, db-nets subsume  $\nu$ -PNs [30], and become expressively equivalent to  $\nu$ -PNs when: (i) there is only one unary color assigning places to the only one unordered countably infinite data type, (ii) the data logic layer is empty. With such a restrictions, the only modeling construct not natively provided by  $\nu$ -PNs is that of arbitrary external input, which can be however simulated using  $\nu$ -PNs by following the strategy defined in [5]. Second, db-nets are expressively equivalent to recently introduced formal models for data-centric business processes, like DCDSs [9] and DMSs [5]. To transform those models into a db-net, it is sufficient to realize a control layer that simulates the application of condition-action rules. The translation of a db-net into those models, instead, is more convoluted, but can be attacked by leveraging the technique introduced in [27] to encode  $\nu$ -PNs into DCDSs. Since DCDSs and DMSs are expressively equivalent to the richest models for business artifacts, such a correspondence paves the way towards the study of CPN-based business artifacts, making approaches like that of [25] truly data-aware.

However, we also stress that db-nets go beyond the aforementioned approaches, since they conceptually componentize the different aspects of a dynamic system, giving first-class citizenships to relations, constraints, queries, database access points in the process, database updates triggered by the process, external inputs, and so on. This opens another interesting line of research, focused on understanding how to exploit db-nets from the conceptual and methodological point of view, as well as on their exploitation to formalize concrete BP management systems like Bizagi, Bonita, and Camunda.

Formal Verification. It is straightforward to see that db-nets, in their full generality, are Turing-complete, and consequently that all standard verification tasks such as reachability, coverability, and model checking are undecidable. However, the fact that the different aspects of a dynamic system are conceptually separated in db-nets make them an ideal model to study in a fine-grained way how such aspects impact on undecidability and complexity of verification tasks, and how should they be controlled to guarantee decidability/tractability. For example, it is known from the literature that the presence of ordered vs. unordered data types, and of (globally) fresh inputs, is intimately connected to the boundaries of decidability for reachability [23]. A similar observation holds for the presence of negation in the queries used to inspect the persistence layer, as well as for the arity of relation schemas contained therein [5]. Interestingly, db-nets do not only provide a comprehensive model to fine-tune all such parameters, but also allow to study how they interact with each other.

Within this space, we consider of particular importance the case where the db-net obeys to the so-called *state-boundedness* property [9,10,27]. Intuitively, in the context of db-nets, this means that the control layer is depth- and width-bounded [30,27], and that the underlying database instance does not simultaneously employ more than a pre-defined number of elements (which however may arbitrarily change over time). Such an assumption still allows for the db-net to visit infinitely many different snapshots along its runs, as no restriction is im-

posed on the size of the type domains, from which external inputs are borrowed. It has been shown that model checking data-aware dynamic systems against first-order variants of  $\mu$ -calculus is indeed decidable, giving a constructive technique to carry out the verification task [9,27]. This opens up another interesting line of investigation on how to check, or guarantee using modeling strategies, that a db-net state-bounded, leveraging recent results [30,10,26,27].

Finally, db-nets paves the way towards the formal analysis of additional properties, which only become relevant when CPNs are combined with relational databases. We mention in particular two families of properties. The first is related to *rollbacks*, so as to check whether it is always (or never) the case that a transition induces a failing action. The second is related to the *true concurrency* present in a db-net, which may contain transitions that appear to be concurrent by considering the control layer in isolation, but have instead to be sequenced due to the interplay with the persistence layer (and its constraints).

Simulation and Benchmarking. Since the control layer of db-nets is grounded on CPNs, all simulation techniques developed for CPNs can be seamlessly lifted to our setting. The result of a db-net simulation produces, as a by-product, a final, database instance, populated through the execution of the control layer. On the one hand, this database instance may be scaled up at one's pleasure, by just changing the simulation parameters. On the other hand, the obtained database instance implicitly reflects the footprint of the control layer, which, e.g., inserts data in a certain order. This makes the obtained database instance much more intriguing than one synthetically generated without considering how data are generated over time. In this light, simulation of db-nets has the potential of providing novel insights into the problem of data benchmarking [22], especially in the context of data preparation for process mining [2,12].

# References

- van der Aalst, W.M.P.: Verification of workflow nets. In: Proc. of ICATPN. LNCS, vol. 1248, pp. 407–426. Springer (1997)
- van der Aalst, W.M.P.: Process cubes: Slicing, dicing, rolling up and drilling down event data for process mining. In: Proc. of AP-BPM. LNCS, vol. 159, pp. 1–22. Springer (2013)
- van der Aalst, W.M.P., Stahl, C.: Modeling Business Processes A Petri Net-Oriented Approach. Cooperative Information Systems series, MIT Press (2011)
- van der Aalst, W.M.P., Weske, M., Grünbauer, D.: Case handling: A new paradigm for business process support. Data and Knowledge Engineering 53(2), 129–162 (2005)
- Abdulla, P.A., Aiswarya, C., Atig, M.F., Montali, M., Rezine, O.: Recency-bounded verification of dynamic database-driven systems. In: Proc. of PODS. ACM Press (2016)
- Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison Wesley (1995)
- Abiteboul, S., Segoufin, L., Vianu, V.: Static analysis of active XML systems. ACM Trans. Database Syst. 34(4) (2009)

- Badouel, E., Hélouët, L., Morvan, C.: Petri nets with semi-structured data. In: Proc. of PN. LNCS, Springer (2015)
- Bagheri Hariri, B., Calvanese, D., De Giacomo, G., Deutsch, A., Montali, M.: Verification of relational data-centric dynamic systems with external services. In: Proc. of PODS. pp. 163–174. ACM (2013)
- Bagheri Hariri, B., Calvanese, D., Deutsch, A., Montali, M.: State boundedness in data-aware dynamic systems. In: Proc. of KR (2014)
- 11. Calvanese, D., De Giacomo, G., Montali, M.: Foundations of data aware process analysis: A database theory perspective. In: Proc. of PODS (2013)
- Calvanese, D., Montali, M., Syamsiyah, A., Aalst, W.M.P.: Ontology-driven extraction of event logs from relational databases. In: Proc. of BPI. LNCS, vol. 256 (2015)
- Cohn, D., Hull, R.: Business artifacts: A data-centric approach to modeling business operations and processes. IEEE Data Eng. Bull. 32(3), 3–9 (2009)
- Damaggio, E., Hull, R., Vaculín, R.: On the equivalence of incremental and fixpoint semantics for business artifacts with guard-stage-milestone lifecycles. Information Systems 38(4), 561–584 (2013)
- De Masellis, R., Di Francescomarino, C., Ghidini, C., Montali, M., Tessaris, S.: Raw-sys: a practical framework for data-aware business process verification. Tech. Rep. KRDB16-1, Free University of Bozen-Bolzano (2016)
- Deutsch, A., Hull, R., Patrizi, F., Vianu, V.: Automatic verification of data-centric business processes. In: Proc. of ICDT. pp. 252–267 (2009)
- Dumas, M.: On the convergence of data and process engineering. In: Proc. of ABDIS. LNCS, vol. 6909, pp. 19–26. Springer (2011)
- Hidders, J., Kwasnikowska, N., Sroka, J., Tyszkiewicz, J., Van den Bussche, J.: Dfl: A dataflow language based on petri nets and nested relational calculus. Inf. Syst. 33(3), 261–284 (2008)
- 19. Hull, R.: Artifact-centric business process models: Brief survey of research results and challenges. In: Proc. of ODBASE. pp. 1152–1163 (2008)
- Jensen, K., Kristensen, L.M.: Coloured Petri Nets Modelling and Validation of Concurrent Systems. Springer (2009)
- Künzle, V., Weber, B., Reichert, M.: Object-aware business processes: Fundamental requirements and their support in existing approaches. Int. J. of Information System Modeling and Design 2(2), 19–46 (2011)
- Lanti, D., Rezk, M., Xiao, G., Calvanese, D.: The NPD benchmark: Reality check for OBDA systems. In: Proc. of EDBT. pp. 617–628. OpenProceedings.org (2015)
- Lasota, S.: Decidability border for petri nets with data: WQO dichotomy conjecture. In: Proc. of PN. LNCS, vol. 9698, pp. 20–36. Springer (2016)
- 24. Libkin, L.: Elements of Finite Model Theory, LNCS, vol. 7360, chap. Fixed Point Logics and Complexity Classes. Springer (2004)
- Lohmann, N.: Compliance by design for artifact-centric business processes. Inf. Syst. 38(4), 606–618 (2013)
- 26. Montali, M., Calvanese, D.: Soundness of data-aware, case-centric processes. Int. Journal on Software Tools for Technology Transfer (2016)
- Montali, M., Rivkin, A.: Model checking petri nets with names using data-centric dynamic systems. Formal Aspects of Computing pp. 1–27 (2016)
- Reichert, M.: Process and data: Two sides of the same coin? In: Proc. of OTM. pp. 2–19 (2012)
- Richardson, C.: Warning: Don't assume your business processes use master data. In: Proc. of BPM. LNCS, vol. 6336, pp. 11–12. Springer (2010)

- 30. Rosa-Velardo, F., de Frutos-Escrig, D.: Decidability and complexity of petri nets with unordered data. Theor. Comput. Sci. 412(34), 4439–4451 (2011)
- Russell, N., ter Hofstede, A.H.M., Edmond, D., van der Aalst, W.M.P.: Workflow data patterns: Identification, representation and tool support. In: Proc. of ER. LNCS, vol. 3716, pp. 353–368. Springer (2005)
- 32. Triebel, M., Sürmeli, J.: Homogeneous equations of algebraic petri nets. In: Proc. of CONCUR. pp. 1–14. LNCS, Springer (2016)
- Vianu, V.: Automatic verification of database-driven systems: a new frontier. In: Proc. of ICDT. pp. 1–13 (2009)