

From DB-nets to Coloured Petri Nets with Priorities (Extended Version)

Marco Montali and Andrey Rivkin

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

Abstract. The recently introduced formalism of DB-nets has brought in a new conceptual way of modelling complex dynamic systems that equally account for the process and data dimensions, considering local data as well as persistent, transactional data. DB-nets combine a coloured variant of Petri nets with name creation and management (which we call ν -CPN), with a relational database. The integration of these two components is realized by equipping the net with special “view” places that query the database and expose the resulting answers to the net, with actions that allow transitions to update the content of the database, and with special arcs capturing compensation in case of transaction failure. In this work, we study whether this sophisticated model can be encoded back into ν -CPNs. In particular, we show that the meaningful fragment of DB-nets where database queries are expressed using unions of conjunctive queries with inequalities can be faithfully encoded into ν -CPNs with transition priorities. This allows us to directly exploit state-of-the-art technologies such as CPN Tools to simulate and analyse this relevant class of DB-nets. (*Topics covered: Higher-level net models, Relationships between Petri nets and other approaches*)

1 Introduction

During the last decade, the Business Process Management (BPM) community has gradually lifted its attention from process models mainly focusing on the flow of activities to multi-perspective models that also account for the interplay between the process and the data perspective [5,12,3]. In particular, several variants of high-level Petri nets have been adopted to capture meaningful integrated models for processes and data, at the same time retaining the possibility of analysing the resulting state space (see, e.g., [10,15,7,11]).

In this spectrum, the recently introduced formalism of DB-nets [11] has brought in a new conceptual way of modelling complex dynamic systems that equally account for the process and data dimensions, considering local data as well as persistent, transactional data. On the one hand, a DB-net adopts a standard relational database with constraints to store persistent data. The database can be queried through SQL/first-order queries, and updated via actions in a transactional way (that is, committing the update only if the resulting




Fig. 1. The conceptual components of DB-nets

database satisfies all intended constraints). On the other hand, a DB-net employs a coloured variant of a Petri net with name creation and management [14] to capture the process control-flow, the injection of (possibly fresh) data such as the creation of new case identifiers [10], and tuples of typed data locally carried out by tokens. This model, which we call ν -CPN, can be seen as a fragment of standard Coloured Petri nets [6] with pattern matching on inscriptions, infinite colour domains, boolean guards, and a very limited use of SML to account for fresh data injection. This also means that ν -CPNs can be seamlessly modelled, simulated, and analysed using state-of-the-art tools such as CPN Tools.

The integration of these two components is realized in a DB-net by extending the ν -CPN with three novel constructs: (i) *view places*, special places that query the database and expose the resulting answers as coloured tokens that can be inspected but not directly consumed; (ii) *action bindings*, linking transitions to database updates by mapping inscription variables to action parameters; (iii) *rollback transition-place arcs*, capturing the emission of tokens in case a fired transition induces a failing database update, and in turn supporting the enablement of compensation transitions. All conceptual components used in the DB-net model are depicted in Figure 1. Notably, DB-nets have been employed to formalize application integration patterns [13].

In this work, we study whether this sophisticated model can be *encoded back into ν -CPNs*, with a twofold intention. On the foundational side, we aim at understanding whether the process-data integration realized in DB-nets adds expressiveness to ν -CPNs, or it is instead conceptual, syntactic sugar. On the practical side, the existence of an encoding would allow us to directly exploit state-of-the-art tools such as CPN Tools towards simulation and analysis of DB-nets. In the case of CPN Tools, this is the only way possible when it comes to state space construction, given the fact that this feature cannot be refined through the third-party extension mechanism offered by the framework.

Specifically, we constructively show through a behavior-preserving translation mechanism that this encoding is indeed possible for a large and meaningful class of DB-nets, provided that the obtained ν -CPN is equipped with *transition priorities* [16] (a feature that is supported by virtually all CPN frameworks, including CPN Tools). Such class corresponds to DB-nets where the database

is equipped with key, foreign key, and domain constraints, and where the view places query the database using unions of conjunctive queries (UCQs) with inequalities. Such query language corresponds to the widely adopted fragment of SQL consisting of select-project-join queries with filters [1].

2 The DB-net Formal Model

In this section, we briefly present the key concepts and notions used for defining DB-nets. Conceptually, a DB-net is composed of three layers (cf. Figure 1) 1) *persistence layer*, capturing a full-fledged relational database with constraints, and used to store background data, and data that are persistent across cases; 2) *control layer*, employing a variant of CPNs to capture the process control-flow, case data, and possibly the resources involved in the process execution; 3) *data logic layer*, interconnecting in the persistence and the control layer.

Definition 1. A db-net is a tuple $\langle \mathcal{D}, \mathcal{P}, \mathcal{L}, \mathcal{N} \rangle$, where: (i) $\mathcal{D}$ is a type domain; (ii) $\mathcal{P}$ is a $\mathcal{D}$ -typed persistence layer; (iii) $\mathcal{L}$ is a $\mathcal{D}$ -typed data logic layer over $\mathcal{P}$; (iv) $\mathcal{N}$ is a $\mathcal{D}$ -typed control layer over $\mathcal{L}$.

We next formalize the framework layer by layer.

Persistence layer. A type domain $\mathcal{D}$ is a finite set of pairwise disjoint data types $\mathcal{D} = \langle \Delta_{\mathcal{D}}, \Gamma_{\mathcal{D}} \rangle$, where $\Delta_{\mathcal{D}}$ is a value domain, and $\Gamma_{\mathcal{D}}$ is a finite set of predicate symbols. Examples of data types are: (i) **string** : $\langle \mathbb{S}, \{=_s\} \rangle$, strings with the equality predicate; (ii) **real** : $\langle \mathbb{R}, \{=_r, <_r\} \rangle$, reals with the usual comparison operators; (iii) **int** : $\langle \mathbb{Z}, \{=_int, <_int, succ\} \rangle$, integers with the usual comparison operators, as well as the successor predicate.

A $\mathcal{D}$ -typed database schema $\mathcal{R}$ is a finite set of $\mathcal{D}$ -typed relation schemas $R(\mathcal{D}_1, \dots, \mathcal{D}_n)$, where $\mathcal{D}_i$ indicates the data type associated to an i -th component of R . A $\mathcal{D}$ -typed database instance $\mathcal{I}$ over $\mathcal{R}$ is a finite set of facts of the form $R(\mathbf{o}_1, \dots, \mathbf{o}_n)$, such that $R(\mathcal{D}_1, \dots, \mathcal{D}_n) \in \mathcal{R}$ and $\mathbf{o}_i \in \Delta_{\mathcal{D}_i}$, for $i \in \{1, \dots, n\}$. Given a type $\mathcal{D} \in \mathcal{D}$, the $\mathcal{D}$ -active domain of $\mathcal{I}$, is the set of $Adom_{\mathcal{D}}(\mathcal{I}) = \{\mathbf{o} \in \Delta_{\mathcal{D}} \mid \mathbf{o} \text{ occurs in } \mathcal{I}\}$.

Given a type domain $\mathcal{D}$, we fix a countably infinite set $\mathcal{V}_{\mathcal{D}}$ of typed variables with a variable typing function $\mathbf{type} : \mathcal{V}_{\mathcal{D}} \rightarrow \mathcal{D}$. As a query language, we adopt standard first-order logic (FOL) extended with data types under the active-domain semantics [8], that is, the evaluation of quantifiers only depends on the values explicitly appearing in the database instance over which they are applied. This can be seen as the FOL representation of SQL queries. A (well-typed) $\mathbf{FO}(\mathcal{D})$ query Q over a $\mathcal{D}$ -typed database schema $\mathcal{R}$ has the form $\{\vec{x} \mid \varphi(\vec{x})\}$, where $\vec{x}$ is the tuple of answer variables of Q , and φ is a FO formula, with $\vec{x}$ as free variables, over predicates in $\cup_{\mathcal{D} \in \mathcal{D}} \Gamma_{\mathcal{D}}$ and relation schemas in $\mathcal{R}$, whose variables and constants are correctly typed. We use $Q(\vec{x})$ to make the answer variables $\vec{x}$ of Q explicit, and denote the set of such variables as $Free(Q)$. When $Free(Q) = \emptyset$, we call Q a boolean query.

A substitution for a set $X = \{x_1, \dots, x_n\}$ of typed variables, is a function $\theta : X \rightarrow \Delta_{\mathcal{D}}$, such that $\theta(x) \in \Delta_{\mathbf{type}(x)}$, for every $x \in X$. A substitution θ for




Fig. 2. The persistence layer for the online shopping scenario

a $\text{FO}(\mathcal{D})$ query Q is a substitution for the free variables of Q . We denote by $Q\theta$ the boolean query obtained from Q by replacing each occurrence of a free variable $x \in \text{Free}(Q)$ with the value $\theta(x)$. Given a $\mathcal{D}$ -typed database schema $\mathcal{R}$, a $\mathcal{D}$ -typed instance $\mathcal{I}$ over $\mathcal{R}$, and a $\text{FO}(\mathcal{D})$ query Q over $\mathcal{R}$, the set of *answers* to Q in $\mathcal{I}$ is defined as the set $\text{ans}(Q, \mathcal{I}) = \{\theta : \text{Free}(Q) \rightarrow \text{Adom}_{\mathcal{D}}(\mathcal{I}) \mid \mathcal{I}, \theta \models Q\}$ of substitutions for Q , where $\models$ denotes standard FO entailment (i.e., we use *active-domain semantics*). We denote by $\text{LIVE}_{\mathcal{D}}(x)$ the unary query returning all the objects of type $\mathcal{D}$ that occur in the active domain (writing such a query is straightforward). When Q is boolean, we write $\text{ans}(Q, \mathcal{I}) \equiv \text{true}$ if $\text{ans}(Q, \mathcal{I})$ consists only of the empty substitution (denoted $\langle \rangle$), and $\text{ans}(Q, \mathcal{I}) \equiv \text{false}$ if $\text{ans}(Q, \mathcal{I}) = \emptyset$. Boolean queries are also used to express *constraints* over $\mathcal{R}$. We introduce explicitly two common types of constraints: given relations R/n and S/m , and two index-sets N and M such that $1 \leq i \leq n$ for every $i \in N$, and $1 \leq j \leq m$ for every $j \in M$, we fix the following notation: (i) $\text{PK}(R) = N$ expresses that the projection $R[N]$ of R on N is a primary key for R ; (ii) $R[N] \subseteq S[M]$ expresses that the projection $R[N]$ of R on N refers the projection $S[M]$ of S on M , which has to be a key for S . Both kinds of constraints are obviously expressible as suitable queries [1].

Definition 2. A $\mathcal{D}$ -typed persistence layer is a pair $\langle \mathcal{R}, \mathcal{E} \rangle$ where: (i) $\mathcal{R}$ is a $\mathcal{D}$ -typed database schema; (ii) $\mathcal{E}$ is a finite set $\{\Phi_1, \dots, \Phi_k\}$ of boolean $\text{FO}(\mathcal{D})$ queries over $\mathcal{R}$, modelling constraints over $\mathcal{R}$.

We say that a $\mathcal{D}$ -typed database instance $\mathcal{I}$ *complies with* $\mathcal{P}$, if $\mathcal{I}$ is defined over $\mathcal{R}$ and satisfies all constraints in $\mathcal{E}$.

Example 1. Let us consider a simplified shopping process used by an e-commerce website. Specifically, we are interested in a simplified scenario in which an already registered user logs in the website and immediately proceeds with selecting products. While products can be selected and added to the shopping cart, the user can occasionally choose a monthly bonus that may be applied when concluding a purchase. We restrict this scenario only by considering cases in which each user ends up buying at least one product.

The persistence layer $\mathcal{P} = \langle \mathcal{R}, \mathcal{E} \rangle$ of this scenario comprises four relation schemas (cf. Figure 2): $\text{User}(\text{int}, \text{string})$ lists registered users together with their credit card data, $\text{WithBonus}(\text{int}, \text{string})$ indicates users that have bonuses, $\text{Product}(\text{string})$ indexes product types offered by the website and $\text{InWarehouse}(\text{int}, \text{string}, \text{real})$ models products (together with their costs) stored in the warehouse. Note the constraints between these schemas. For example, in order to show that users cannot have more than one bonus at a time, we introduce a foreign key constraint between WithBonus and User that is denoted as $\text{WithBonus}[\{1\}] \subseteq \text{User}[\{1\}]$ and formalized in FO

logic as: $\forall uid, bt. WithBonus(uid, bt) \rightarrow \exists card. User(uid, card)$. Another constraint limits the bonus type values in *WithBonus* and can be expressed as $\forall uid, bt. WithBonus(uid, bt) \rightarrow bt = 50\% \vee bt = 15eur \vee bt = \text{extra_item}$. $\square$

Data logic layer. The data logic layer allows one to *extract* data from the database instance using queries as well as to *update* the database instance by adding and deleting possibly multiple facts at once. The updates follow the *transactional* semantics: if a new database instance obtained after some update is still compliant with the persistence layer, the update is *committed*; otherwise it is *rolled back*. Such updates are realized in parametric atomic actions, resembling ADL actions in planning [4], and consist of fact templates – expressions that, once instantiated, assert which facts will be added to and deleted from the database. Specifically, given a typed relation $R(\mathcal{D}_1, \dots, \mathcal{D}_n) \in \mathcal{R}$, an *R*-fact template over $\vec{p}$ has the form $R(y_1, \dots, y_n)$, such that for every $i \in \{1, \dots, n\}$, y_i is either a value $o \in \Delta_{\mathcal{D}_i}$, or a variable $x \in \vec{p}$ with $\text{type}(x) = \mathcal{D}_i$.

A (*parameterized*) *action* over a $\mathfrak{D}$ -typed persistence layer $\langle \mathcal{R}, \mathcal{E} \rangle$ is a tuple $\langle n, \vec{p}, F^+, F^- \rangle$, where: (i) n is the *action name*; (ii) $\vec{p}$ is a tuple of pairwise distinct variables from $\mathcal{V}_{\mathfrak{D}}$, denoting the *action (formal) parameters*; (iii) F^+ and F^- respectively represent a finite set of $\mathcal{R}$ -fact templates (i.e., some *R*-fact templates for some $R \in \mathcal{R}$) over $\vec{p}$, to be *added* to and *deleted* from the current database instance. To access the different components of an action α , we use a dot notation: $\alpha.\text{name} = n$, $\alpha.\text{params} = \vec{p}$, $\alpha.\text{add} = F^+$, and $\alpha.\text{del} = F^-$. Given an action α and a (parameter) substitution θ for $\alpha.\text{params}$, we call *action instance* $\alpha\theta$ the (ground) action resulting by substituting parameters of α with corresponding values from θ . Then, given a $\mathfrak{D}$ -typed database instance $\mathcal{I}$ compliant with $\mathfrak{D}$, the *application* of $\alpha\theta$ on $\mathcal{I}$, written $\text{apply}(\alpha\theta, \mathcal{I})$, is a database instance over $\mathcal{R}$ obtained as $(\mathcal{I} \setminus F_{\alpha\theta}^-) \cup F_{\alpha\theta}^+$, where: (i) $F_{\alpha\theta}^- = \bigcup_{R(\vec{y}) \in \alpha.\text{del}} R(\vec{y})\theta$; (ii) $F_{\alpha\theta}^+ = \bigcup_{R(\vec{y}) \in \alpha.\text{add}} R(\vec{y})\theta$. If $\text{apply}(\alpha\theta, \mathcal{I})$ complies with $\mathcal{P}$, $\alpha\theta$ can be *successfully applied* to $\mathcal{I}$. Note that, in order to avoid situations in which the same fact is asserted to be added and deleted, we prioritize deletions over additions.

Definition 3. 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 $\text{FO}(\mathfrak{D})$ 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 $\mathcal{P}$. To inspect the persistence layer, we use the following queries:

- $\mathbf{Q}_{\text{products}}(pid, n, c) :- \text{Product}(n) \wedge \text{InWarehouse}(pid, n, c) \wedge c \neq \text{null}$, to extract products available in the warehouse and whose price is not *null* (those without prices can be undergoing the stock-taking process);
- $\mathbf{Q}_{\text{users}}(uid) :- \exists card. \text{User}(id, card)$, to get registered users;
- $\mathbf{Q}_{\text{wbonus}}(uid, bt', u) :- \text{WithBonus}(uid, bt', u)$, to inspect all users with bonuses.

In addition, $\mathcal{L}$ provides key functionalities for organizing the shopping process. Such functionalities are realized through four actions (where, for simplicity, we blur the distinction between an action and its name). To manage bonuses we use two actions *ADDB* and *CHANGE*. The former is used to assign a bonus of type

bt to a user with id uid ($\text{ADDB}\cdot\text{params} = \langle uid, bt \rangle$) and record it into the persistent storage: $\text{ADDB}\cdot\text{add} = \{ \text{WithBonus}(uid, bt) \}$, $\text{ADDB}\cdot\text{del} = \emptyset$. Note that, before logging in, the user may have already a bonus assigned during one of the previous sessions. At will, such a bonus can be changed using the following action: $\text{CHANGE}\cdot\text{add} = \{ \text{WithBonus}(uid, bt') \}$, $\text{CHANGE}\cdot\text{del} = \{ \text{WithBonus}(uid, bt) \}$. In fact, CHANGE realizes an update by first deleting a tuple that is characterized by uid and bt' (the old bonus), and then adding its modified version. We use RESERVE to reserve product pid stored in cart cid for further processing (e.g., the preparation for shipment) by deleting it from the list of available products: $\text{RESERVE}\cdot\text{add} = \emptyset$, $\text{RESERVE}\cdot\text{del} = \{ \text{InWarehouse}(pid, n, c) \}$. At last, we may utilize our monthly bonus (if it has not been yet used) to consider it when paying the order. For that, we use an action called APPLY such that: $\text{APPLY}\cdot\text{add} = \emptyset$, $\text{APPLY}\cdot\text{del} = \{ \text{WithBonus}(uid, bt) \}$. $\square$

Control layer. The control layer employs a fragment of Coloured Petri net to capture the process control flow and a data logic to interact with an underlying persistence layer. 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 $m : A \rightarrow \mathbb{N}$. 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 $k \in \mathbb{N}$, $k \cdot S_1 = \{a^{kn} \mid a^n \in S_1\}$.

We shall call *inscription* a tuple of typed variables (and, possibly, values) and denote the set of all possible inscriptions over set $\mathcal{Y}$ as $\Omega_{\mathcal{Y}}$, and the set of variables appearing inside an inscription $\omega \in \Omega_{\mathcal{Y}}$ as $\text{Vars}(\omega)$ (such notation naturally extends to sets and multisets of inscriptions). In the spirit of CPNs, the control layer assigns to each place a color type, which in turn combines one or more data types from $\mathfrak{D}$. Formally, a $\mathfrak{D}$ -color is $\mathcal{D}_1 \times \dots \times \mathcal{D}_m$, where for each $i \in \{1, \dots, m\}$, we have $\mathcal{D}_i \in \mathfrak{D}$. We denote by Σ the set of all possible $\mathfrak{D}$ -colors. To account for fresh external inputs, we employ the well-known mechanism adopted in ν -Petri nets [14,10] and introduce a countably infinite set $\mathcal{Y}_{\mathfrak{D}}$ of $\mathfrak{D}$ -typed *fresh variables*. To guarantee an unlimited provisioning of fresh values, we impose that for every variable $\nu \in \mathcal{Y}_{\mathfrak{D}}$, we have that $\Delta_{\text{type}(\nu)}$ is countably infinite. Hereinafter, we shall fix a countably infinite set of $\mathfrak{D}$ -typed variable $\mathcal{X}_{\mathfrak{D}} = \mathcal{V}_{\mathfrak{D}} \cup \mathcal{Y}_{\mathfrak{D}}$ as the disjoint union of “normal” variables $\mathcal{V}_{\mathfrak{D}}$ and fresh variables $\mathcal{Y}_{\mathfrak{D}}$.

As we have mentioned before, the control layer can be split into two parts. Let us first define the ν -CPN part that can be seen as an extension of ν -Petri nets with concrete data types, boolean (type-aware) guards and read arcs.

Definition 4. A $\mathfrak{D}$ -typed ν -CPN $\mathcal{N}$ is a tuple $\langle P, T, F_{in}, F_{out}, \text{color} \rangle$, where:

1. P is a finite set of places.
2. $\text{color} : P \rightarrow \Sigma$ is a color type assignment over P mapping each place $p \in P$ to a corresponding $\mathfrak{D}$ -type color.

3. T is a finite set of transitions, such that $T \cap P = \emptyset$.
4. $F_{in} : P \times T \rightarrow \Omega_{\mathcal{V}_{\mathfrak{D}}}^{\oplus}$ is an input flow from P to T assigning multisets of inscriptions (over variables $\mathcal{V}_{\mathfrak{D}}$) to input arcs, s.t. that each of such inscriptions $\langle x_1, \dots, x_m \rangle$ is compatible with each of its input places p , i.e., for every $i \in \{1, \dots, m\}$, we have $\mathbf{type}(x_i) = \mathcal{D}_i$, where $\mathbf{color}(p) = \mathcal{D}_1 \times \dots \times \mathcal{D}_m$.
5. $\mathbf{guard} : T \rightarrow \mathbb{F}_{\mathfrak{D}}$ is a transition guard assignment over T assigning to each transition $t \in T$ a $\mathfrak{D}$ -typed guard φ , s.t.:
 - $\text{InVars}(t) = \{x \in \mathcal{V}_{\mathfrak{D}} \mid \text{there exists } p \in P \text{ such that } x \in \text{Vars}(F_{in}(\langle p, t \rangle))\}$ is the set of all variables occurring on input arc inscriptions of t ;
 - a $\mathfrak{D}$ -typed guard from is a formula (or a quantifier- and relation-free $\text{FO}(\mathfrak{D})$ query) of the form $\varphi ::= \text{true} \mid S(\vec{y}) \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2$, where $S/n \in \Gamma_{\mathfrak{D}}$ and, for $\vec{y} = \langle y_1, \dots, y_n \rangle \subseteq \mathcal{V}_{\mathfrak{D}}$, we have that y_i is either a value $o \in \Delta_{\mathfrak{D}}$, or a variable $x_i \in \mathcal{V}_{\mathfrak{D}}$ with $\mathbf{type}(x_i) = \mathcal{D}$ ($i \in \{1, \dots, n\}$);
 - $\mathbb{F}_{\mathfrak{D}}$ is the set of all possible $\mathfrak{D}$ -typed guards and, with a slight abuse of notation, $\text{Vars}(\varphi)$ is the set of variables occurring in φ .
6. $F_{out} : T \times P \rightarrow \Omega_{\mathcal{X}_{\mathfrak{D}} \cup \Delta_{\mathfrak{D}}}^{\oplus}$ is an output flow from transitions T to places P assigning multisets of inscriptions to output arcs, such that all such inscriptions are compatible with their output places.

According to the diagram in Figure 1, the DB-net control layer can be obtained on top of ν -CPNs by essentially adding three mechanisms that allow the net to interact with the underlying persistent storage: (i) view places, allowing the net to inspect parts of the database using queries; (ii) action binding, linking atomic actions and their parameters to transitions and their inscription variables; (iii) rollback transition-place arcs, enacted when the action application induced by a transition firing violates some database constraint, so as to explicitly account for “error-handling”.

Definition 5. A $\mathfrak{D}$ -typed control layer over a data logic layer $\mathcal{L} = \langle \mathcal{Q}, \mathcal{A} \rangle$ is a tuple $\langle P, T, F_{in}, F_{out}, F_{rb}, \mathbf{color}, \mathbf{query}, \mathbf{guard}, \mathbf{act} \rangle$, where:

1. $\langle P_c, T, F_{in}, F_{out}, \mathbf{color} \rangle$ is a $\mathfrak{D}$ -typed ν -CPN, where P_c is a finite set of control places.
2. $P = P_c \cup P_v$ is a finite set of places, where P_v are view places (decorated as $\textcircled{\bullet}$ and connected to transitions with special read arcs).
3. $\mathbf{query} : P_v \rightarrow \mathcal{Q}$ is a query assignment mapping each view place $p \in P_v$ with $\mathbf{color}(p) = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$ to a query $Q(x_1, \dots, x_n)$ from $\mathcal{Q}$, s.t. the color of p component-wise matches with the types of the free variables in Q : for each $i \in \{1, \dots, n\}$, we have $\mathcal{D}_i = \mathbf{type}(x_i)$.
4. $\mathbf{act} : T \rightarrow \mathcal{A} \times \Omega_{\mathcal{X}_{\mathfrak{D}} \cup \Delta_{\mathfrak{D}}}$ is a partial function assigning transitions in T to actions in $\mathcal{A}$, where $\mathbf{act}(t)$ maps t to an action $\alpha \in \mathcal{A}$ together with a (binding) inscription $\langle y_1, \dots, y_m \rangle$, s.t. if $\alpha.\mathbf{params} = \langle z_1, \dots, z_m \rangle$ and, for each $i \in \{1, \dots, m\}$, we have $\mathbf{type}(y_i) = \mathbf{type}(z_i)$ if y_i is a variable from $\mathcal{X}_{\mathfrak{D}}$, or $y_i \in \Delta_{\mathbf{type}(z_i)}$ if y_i is a value from $\Delta_{\mathfrak{D}}$.
5. F_{rb} is an output flow from T to P_c called rollback flow (we shall refer to F_{out} as normal output flow).




Fig. 3. The control of a DB-net for online shopping. Here, ν_c is a fresh input variable corresponding to a newly created cart, whereas $dest$ is an arbitrary input variable representing a destination address. The rollback output arc (corresponds to the rollback flow) is in red and decorated with an “x”.

Figure 3 shows the control layer of the shopping cart example. The queries specified in Example 2 are assigned to the corresponding view places: $\text{query}(\text{Products}) := Q_{\text{products}}$, $\text{query}(\text{Users}) := Q_{\text{users}}$ and $\text{query}(\text{Bonus Holders}) := Q_{\text{wbonus}}$. The actions (with their formal parameters) assigned to transitions via **act** graphically appear in grey transition boxes.

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. Due to space limitations, we refer to the definition of the formal semantics studied in [11]. We thus assume that the execution semantics of both ν -CPNs and DB-nets can be captured with a possibly infinite-state labeled transition system (LTS) that accounts for all possible executions starting from their initial markings. While transitions in such LTSs model the effect of firing nets under given bindings, their state representations slightly differ. Namely, in the case of ν -CPNs we have markings (like, for example, in coloured Petri nets [6]), while in the case of DB-nets one also has to take into account database states. W.l.o.g., we shall use $\Gamma_{\mathcal{N}}^{M_0} = \langle \mathcal{M}, M_0, \rightarrow, L \rangle$ to specify an LTS for a ν -CPN $\mathcal{N}$ with initial marking M_0 and $\Gamma_{\mathcal{B}}^{s_0} = \langle S, s_0, \rightarrow, L \rangle$ to specify an LTS for a DB-net $\mathcal{B}$ with initial snapshot $s_0 = \langle \mathcal{I}_0, m_0 \rangle$, where $\mathcal{I}_0$ is the initial database instance and m_0 is the initial marking of the control layer.

3 Translation

We are now ready to describe the translation from DB-nets to ν -CPNs with priorities (we assume the reader is familiar with transition priorities). Recall that this is not just of theoretical interest, but has also practical implications. In [13], we have presented a prototypical implementation of DB-nets in CPN Tools that, using Access/CPN and Comms/CPN, allow to model and simulate DB-nets. However, we realized that CPN Tools would not correctly generate




Fig. 4. A generic DB-net transition accessing multiple view places

the state space of the DB-net at hand. This is due to the fact that the CPN Tools state space construction module does not consider third-party extensions which, in our setting, implies that the content of the view places is not properly recomputed after each transition firing.

The first challenge to overcome is how the database schema is represented in the target net. To this aim, we introduce special *relation places* that copy corresponding database relations by mirroring their signature to the type definitions of places.¹ In this light, database instances will correspond to relation place markings, where tokens are nothing but tuples. All other DB-net elements (for example, bindings for fresh variables, action execution) require actual computation that happens when a transition fires. Intuitively, every DB-net transition T is represented using the following four phases:

1. Collect bindings and compute the content of view places adjacent to T .
2. If there is an action assigned to T , execute it. We employ auxiliary boolean places that control whether an update has actually happened (that is, a token representing a tuple has been removed from or added to a relation place).
3. Check the satisfaction of integrity constraints.
4. Finish the computation and generate a new marking.
 - (a) If all constraints are satisfied, empty the auxiliary boolean places used in 2), release the lock, and populate the postset of T .
 - (b) If some constraint is violated, roll-back the effects. This is done in reverse order w.r.t. phase 2), applying or skipping a reverse update depending on how the values in the special places. After this, the relation places have the content they had before the action was applied. Then, one releases the lock and pushes the special postset corresponding to the roll-back arc (if any) attached to T .

To realize the execution of an original DB-net transition, all the four phases are executed uninterruptedly (under lock). In the reminder of the section we formalizing the phases discussed above.

A generic DB-net $\mathcal{B}_\tau$ that we use to demonstrate the translation is represented in Figure 4. Here, we assume that T contains enough of tokens assigned by its input flow and its eventual firing is subject to the $G(\vec{y})$ guard evaluation. $\vec{y}$, in turn, is bound to values from $\vec{z}$ and from $m \in \mathbb{N}$ ordered view places, where each view place V_i has a query Q_{v_i} assigned to it. The ν -CPN $\mathcal{N}_\tau$ representing $\mathcal{B}_\tau$ is depicted in Figure 5. To facilitate the translation, we make three working hypothesis. First, we assume that the relational schema is equipped only with

¹ Relation places do not differ from the normal ν -CPN places. We use the different name in order to conceptually distinguish their origin.




Fig. 5. Overall ν -CPN encoding of the DB-net transition shown in Figure 4. Blue clouds stand for subnets that are expanded next, and $\vec{x}$ is a shortcut for the tuple consisting of $\vec{x}_1 \dots, \vec{x}_m$. Elements within the gray rectangle are local to the transition, whereas external elements are shared at the level of the whole net.

three types of constraints: primary keys, foreign keys and domain constraints. Second, for ease of presentation, we assume that the resulting ν -CPN model can deal with DB-nets external variables. This assumption, however, is correct from the practical point of view as it has been already shown before that a preliminary implementation of DB-nets in CPN Tools [13] provides functionality necessary for computing bindings for external variables. Third, we naturally extend the notion of ν -CPN with read arcs.




Fig. 6. Fragments of the process and persistence layers of the DB-net in Figure 3




Fig. 7. A ν -CPN representation of the DB-net in Figure 7

3.1 Computing Views Using CPN Places

We start by describing how the view computation should work using only ν -CPN places. Let us consider as an example a subnet $\mathcal{B}_{tr}$ of the DB-net present in Figure 3 that models only the selection of available products. To access products that are available in the warehouse and that have prices assigned to them, we need to run a query $Q_{\text{products}}(pid, n, c) = Product(n) \wedge InWarehouse(pid, n, c) \wedge c \neq \text{null}$. Interestingly, such a query can be formulated directly using standard elements of ν -CPNs. Indeed, we may transfer the DB-net in Figure 6 into a ν -CPN $\mathcal{N}_{tr}$ in Figure 7 representing the project selection step. As one can see, the relations of $\mathcal{B}_{tr}$ have been copied to the same-named relation places, when Q_{products} is treated as follows: $\mathcal{N}_{tr}$ accesses relation places with read-arcs (that have relation attributes as their inscriptions) so as to realize the projection, while the filter (i.e., $c \neq \text{null}$) is basically plugged into the guard of **Add Product**. The result of the query is then propagated into the post-set of **Add Product** using the free variables of Q_{products} (i.e., pid , n and c) in the arc inscriptions.

However, one may see that not every query can be handled when only using standard ν -CPN elements. Assume a query $Q_{\neg\text{available}}(n) = Product(n) \wedge \neg \exists pid, c. InWarehouse(pid, n, c)$ that lists products not available in the warehouse. In order to represent $Q_{\neg\text{available}}$ in a ν -CPN, one would need to extend the net with constructs allowing to fire a transition only if a certain element does not exists in a place incident to it. Thus, we restrict ourselves to the union of conjunctive queries with negative filters (or atomic negations) (UCQFs $^\neq$), that is $F0(\mathcal{D})\mathcal{D}$ queries of the form $\bigvee_{i=1}^n \exists \vec{y}_i. conj_i(\vec{x})$, where $conj_i(\vec{x})$ is also a $F0(\mathcal{D})\mathcal{D}$ query that is a conjunction of relations $R(\vec{z})$, predicates $P(\vec{y})$ and their negations $\neg P(\vec{y})$. Henceforth, we use $Q^{UCQF^\neq}$ to define a UCQF $^\neq$ subset of $\mathcal{Q}$. In SQL, a conjunctive query is a query representable with a SELECT-FROM-WHERE expression. As it has been already shown, the filter conditions (of the UCQFs $^\neq$ attached to view places) can be modeled using transition guards.

In case of multiple view places attached to one transition, we construct a net that computes them in a sequential manner. One may see the computation process as a pipeline. Whenever a transition that corresponds to a certain view




Fig. 8. Expansion of the binding net from Figure 5

place is enabled, it fires and generated tokens that represent one of the tuples of the view. Then, acquired tokens are transferred to the next transition using variables in the arc inscriptions. The computation continues until the last view. After that, the results of all the computations are transferred to the corresponding places, following the topology (i.e., the organization of arcs defined by the flow relations) of the original DB-net. Note that the order in which views are computed has to be the same as the one defined for $\mathcal{B}_\tau$.

A ν -CPN in Figure 8 shows how bindings and view places are computed in the case of the generic DB-net $\mathcal{B}_\tau$. The computation process per view V_i is realized by a transition called *Compute V_i* and analogous to the one explained before: we read necessary data from relation places, representing relations used in $\mathbf{Q}_{V_i}$, and filter these data by means of $F_{V_i}(\vec{y})$. Note that variables on every read-arc adjacent to *Compute V_i* represent attributes of some relation R . The intermediate result of the view computation is then stored in a place called *V_i Computed*. As one can see from Figure 8, all the intermediate results are accumulated along the computation cycle. Moreover, we carry data provided with input variables of $\mathbf{T}$ so as to check the validity of the guard G (see Figure 5). This is done using prioritized transition T_{cond} . If the guard is not satisfied, one has to reset the computation process by returning tokens that have been consumed at the beginning of the view computation (that is, tokens that have been assigned to z). We resolve this issue by introducing an auxiliary transition called T_{cancel} that may fire only when the guard has been evaluated to false. Scheduling between T_{cond} and T_{cancel} is managed by means of two priority labels **P_High** and **P_Low** (where **P_High** > **P_Low**) respectively assigned to them.

3.2 Modeling RDBMS updates in CPNs

We now show how database updates exploited by DB-nets could be represented using regular coloured Petri nets. We recall, that actions assigned to DB-net transitions support addition and deletion of $\mathcal{R}$ -fact, which should preserve the set semantics adopted by the persistence layer.




Fig. 9. A subnet of the DB-net in Figure 3 describing the bonus change step

In Figure 9 we consider a DB-net describing the bonus change step of the online shopping process. Here, for ease of presentation, instead of considering a view place for bonus holders, we use a regular (control) place that stores the same kind of data.





Fig. 10. The CPN representation of the DB-net in Figure 9, where ξ , for ease of reading, denotes the tuple $\langle uid, cid, bt, bt' \rangle$

The translation of DB-net-like database updates into ν -CPNs is conceptually similar to the representation of the view computation process: DB-net actions must be performed sequentially within a critical section that can be entered whenever a special write lock is available. For preserving the set semantics over every relation place, we use prioritized transitions so as to check whether a tuple to be added or deleted already exists in the relation place. Specifically, for each tuple we would introduce two transitions, one with a higher priority and another with a lower priority, and an auxiliary (*no-op*) boolean place. The first transition can fire if the tuple is in the corresponding relation place, while the second one would fire otherwise. Both transitions are adopted to deal with additions and deletions. In case of additions, the highly prioritized transition would not add the tuple, while the one with the lower priority would do otherwise. To deal with deletions, we mirror the previous case: if the tuple exists, then one can safely remove it; otherwise, one proceeds without changes. Upon firing of any of these transitions, the auxiliary place receives a boolean token. If the value of the token is *true*, then it means that the tuple has been successfully added or deleted. In case the database update has not taken place, the token value is going to be *false*. It is important to note that the update execution order of DB-net actions must be also preserved in their ν -CPN representation. That is, for every action α we first delete all the tuples from $\alpha \cdot \text{del}$, and only then add those from $\alpha \cdot \text{add}$.

We incorporate aforementioned modeling guidelines in the ν -CPN depicted in Figure 10. Since CHANGE in $\mathcal{B}_{tr}^\alpha$ contains multiple database updates, the model starts with deleting $\text{WithBonus}(uid, bt)$ from WithBonus . To do so, at first one checks whether the relation place WithBonus contains the tuple we would like to remove. This is done using **ExistsD1** that performs conditional removal of $\text{WithBonus}(uid, bt)$, that is, if there is a token in WithBonus such that bindings of inscriptions on $(D1, \text{ExistsD1})$ and $(\text{WithBonus}, \text{ExistsD1})$ coincide, then **ExistsD1** is enabled, and upon firing consumes the selected token from WithBonus and populates one token with value $\langle \text{true} \rangle$ (the value *true* means that the update has been successfully accomplished) in DoneD1 . Note that **ExistsD1** is always checked first given the higher priority label assigned to it. If the tuple



(a) Expansion of the i -th deletion component in the net of Figure 11 (b) Expansion of the i -th addition component in the net of Figure 11

Fig. 12. Expansion of deletion and addition nets from Figure 11





Fig. 13. Expansion of the check constraint net from Figure 5

the first and the last one could be relatively easy to check during the update phase, assuming that the computation results are accumulated in arc inscriptions analogously to the binding net in Figure 8², the process of managing updates in the presence of arbitrary many foreign key dependencies is quite involved. To manage it correctly we first perform the updates and only then check whether the generated marking represents a database instance that satisfies all the integrity constraints contained in the persistence layer of $\mathcal{B}_\tau$. A ν -CPN representing the check constraint phase is depicted in Figure 13. The net works as follows: it consequently runs small nets for verifying the integrity of constraints and, in case of violation, puts a token in a special place called *ConstrViol*. As soon as there is at least one token in *ConstrViol* place, the big net in Figure 5 terminates the constraint checking process and switches to the phase 4.(b) (that is, runs the undo net) explained in the beginning of this section. For ease of presentation we

² Both primary keys and domain constraints can be violated when a tuple is about to be inserted into a table. Specifically, to guarantee that primary keys are respected, it is enough to check with $ExistsA_i$ whether there is a token in $R_{a,i}$ that has the same primary key value, and, if so, cancel the computation process. In the case of check constraints, one may insert a third transition that has a normal priority and that will be fired whenever one of the values we want to insert is not falling into the allowed range. Firing of this transition will have the same consequences as in the case of primary keys.



(a) Expansion for a primary key constraint C_i indicating that the first k components of relation R with arity $n_r \geq k$ form a key of relation R with arity $n_r \geq m$ reference for R . In the figure, $\vec{y}$ and $\vec{w}$ are tuples containing n_r variables.

(b) Expansion for a foreign key constraint C_i indicating that the last m components of relation S with arity $n_s \geq m$. In the figure, $\vec{y}$ and $\vec{w}$ are tuples containing n_r and n_s variables.

Fig. 14. Expansion of the check net from Figure 13 in the case of key constraints

assume that every relation R/n_r from $\mathcal{R}$ of $\mathcal{B}_\tau$ will have the following look: the first k attributes form a primary key, while the rest of $n_r - k + 1$ attributes can be unconstrained or bounded by domain constraints. Moreover, if R is referencing some other relation S , then among these $n_r - k + 1$ attributes we reserve the last m such that $S[\{1, \dots, m\}] \subseteq R[\{n_r - m, \dots, n_r\}]$.

The constraint checking process starts with verifying that all the updates performed using the net in Figure 11 are satisfying primary key constraints C_i . This is done by sequentially running small check nets in Figure 14(a), where the constraint integrity is verified for some relation R by a pair of prioritized transitions *RepeatedKey* (high priority) and *NoRepeatedKey* (low priority). Note that *RepeatedKey* accesses the content of R with two read arcs and using the guard assigned to it verifies whether there exist two tokens, such that their first k values coincide and in the rest of $n_r - k + 1$ values there is at least one distinct pair of values. The satisfaction of such guard would mean that, essentially, we have inserted a token in R whose primary key values were not unique. Firing of *RepeatedKey* will produce one token in *ConstrViol* and terminate the run of the check constraint net.

The next type of constraints to verify is the foreign key dependency. Analogously to the previous case, we successively run small check nets like the one in Figure 14(b) and in each of them control that R correctly references S (that is, there are no tuples in R that do not depend on any tuple in S). This is realized with two prioritized transitions *FKExists* (high priority) and *FKNotExists* (low priority). The first one, as the name suggests, checks whether the dependency between R and S is preserved for all the tokens in the corresponding relation places. *FKExists* makes use of the guard attached to it that performs pairwise comparison of m last values of a token from R to m first values of a token from S . If the guard is not satisfied, then the dependency relation between R and S has been violated and one fires *FKNotExists* so as to terminate the constraint checking process.




Fig. 15. Expansion of the check net from Figure 13 in the case of a domain constraint C_i indicating that the j -th component of relation R with arity $n_r \geq j$ must contain a value that belongs to $\{c_1, \dots, c_l\}$; In the figure, $\vec{y}$ is a tuple containing n_r variables.

The last series of constraints to be checked is the one of domain constraints. The net in Figure 15 employs two prioritized transitions, *WrongValue* and *NoWrongValue*, to verify whether all the tuples inserted into R had correct values. First, using the guard of *WrongValue*, we check whether there is at least one value that breaks the integrity of the domain constraint C_i of R . If *WrongValue* fires, the process is terminated by putting a token into *ConstrViol*. Otherwise, *NoWrongValue* is executed and the constraint checking process continues.

Now let us show how the computation of all the effects of T is finished and a new marking is generated. If one of the constraints has been violated, we have to roll back all the effects pushed using the net in Figure 11. To do so, we employ the net in Figure 16 that reverts the update process by first canceling all the additions, and then canceling all the deletions. Let us briefly explain how the rollback process is performed for each component net.




Fig. 16. Expansion of the undo net from Figure 5

We start by removing all the tuples that have been successfully added to relation places following the definition of ACT. The net in Figure 17(a) shows how to revert the result of inserting $R_{a,i}$ -fact from F^+ (i.e., a fact of the form $R_{a,i}(\vec{y}_{a,i})$). If a fact has been added, that is, there is a token with value $\langle \text{true} \rangle$ in $DoneA_i$, then the net removes it by firing $DoRevertA_i$. Otherwise, if the fact has not been added, that is, there is a token with value $\langle \text{false} \rangle$ in $DoneA_i$, then the net proceeds without reverting by firing $SkipRevertA_i$. Then, for each $R_{d,i}$ -fact, we go on with adding all the tuples that have been deleted by using the net depicted in Figure 17(b). The update reverting process is analogous to the one dealing with reverted additions, but with only one exception: whenever $DoneD_i$ has a $\langle \text{true} \rangle$ token, then we put the deleted tuple (specified in F^-) back into $R_{d,i}$. Note that every revert deletion or addition net removes a token from a corresponding auxiliary no-op place.

As soon as all the operations of ACT have been undone and all the corresponding tokens have been withdrawn from relation places, the net places a token




Fig. 17. Expansion of revert deletion and addition nets from Figure 16




Fig. 18. Expansion of the consume net from Figure 5

in *DoRollBack* (cf. Figure 5) and allows us to fire a transition called $T_{rollback}$ that implements the generation of the tokens in the postset corresponding to the rollback flow of T .

If the check constraint net's work has not been interrupted an the token was placed in *ConstrOk* (cf. Figure 5), then we proceed with the consume net (cf. Figure 18) that removes all tokens from the auxiliary no-op places and places a token into *DoCommit*. This, in turn, allows $\mathcal{N}_\tau$ to execute T_{commit} that populates tokens in the postset corresponding to the normal flow of T .

3.4 The general translation

In this section we bring together the modeling approaches described in the previous three sections and quickly summarize the translation from DB-nets to ν -CPNs with priorities. Specifically, we show that, given a DB-net, it is possible to build a ν -CPN that is weakly bisimilar to it.

Intuitively, $\mathcal{N}_\tau$ from Figure 5 behaves just like the $\mathcal{B}_\tau$ in Figure 4 and hence LTSs of these two nets are bisimilar [9]. Notice that, in order to correctly represent the behavior of $\mathcal{B}_\tau$, $\mathcal{N}_\tau$ includes many intermediate steps that are, however, not relevant for comparing content of the states and behavior of the nets. For this we are going to resort to a form of bisimulation that allows to “skip” transitions

irrelevant for the behavioral comparison [9]. Specifically, given two transition systems $\Gamma_1 = \langle S_1, s_{01}, \rightarrow_1, L \rangle$ and $\Gamma_2 = \langle S_2, s_{02}, \rightarrow_2, L \rangle$ defined over a set of labels L , we call relation $wb \subseteq S_1 \times S_2$ a *weak bisimulation* between Γ_1 and Γ_2 iff for every pair $\langle p, q \rangle \in wb$ and $a \in L \cup \{\epsilon\}$ the following holds: (1) if $p \xrightarrow{a}_1 p'$, then there exists $q' \in S_2$ such that $q \xrightarrow{a}_2 q'$ and $\langle p', q' \rangle \in wb$; (2) if $q \xrightarrow{a}_2 q'$, then there exists $p' \in S_1$ such that $p \xrightarrow{a}_1 p'$ and $\langle p', q' \rangle \in wb$. Here, $\epsilon \neq a$ is a special silent label and $p \xrightarrow{a}_1 q$ is a weak transition that is defined as follows: i) $p \xrightarrow{\epsilon}_1 q$ iff $p(\xrightarrow{\epsilon})^*_1 q_1 \xrightarrow{a}_1 q_2(\xrightarrow{\epsilon})^*_1 q$; ii) $p \xrightarrow{\epsilon}_1 q$ iff $p(\xrightarrow{\epsilon})^*_1 q$. We use $(\xrightarrow{\epsilon})^*$ to define the reflexive and transitive closure of $\xrightarrow{\epsilon}$. We say that a state $p \in S_1$ is weakly bisimilar to $q \in S_2$, written $p \approx^{wb} q$, if there exists a weak bisimulation wb between Γ_1 and Γ_2 such that $\langle p, q \rangle \in wb$. Finally, Γ_1 is said to be weakly bisimilar to Γ_2 , written $\Gamma_1 \approx^{wb} \Gamma_2$, if $s_{01} \approx^{wb} s_{02}$. Let us now define a theorem that sets up the behavioral correspondence between DB-nets and ν -CPNs.

Theorem 1. *Let $\mathcal{B} = \langle \mathcal{D}, \mathcal{P}, \mathcal{L}, \mathcal{N} \rangle$ be a DB-net with $\mathcal{P} = \langle \mathcal{R}, \emptyset \rangle$, $\mathcal{L} = \langle \mathcal{Q}^{UCQF^\neq}, \mathcal{A} \rangle$ and $\mathcal{N} = \langle P, T, F_{in}, F_{out}, F_{rb}, \text{color}, \text{query}, \text{guard}, \text{act} \rangle$, and s_0 is the initial snapshot. Then, there exists a ν -CPN $\mathcal{N} = \langle P \cup P_{rel} \cup P_{aux}, T \cup T_{aux}, F_{in}, F_{out}, \text{color}, M_0 \rangle$ with (1) a set of relation places P_{rel} acquired from $\mathcal{R}$, (2) two sets P_{aux} and T_{aux} of auxiliary places and transitions (required by the encoding algorithm), and such that $\Gamma_{\mathcal{B}}^{s_0} \approx_{flat}^{wb} \Gamma_{\mathcal{N}}^{M_0}$.*

The proof of the theorem is obtained inductively by modularly considering the encoding defined in Sections 3.1–3.3. Intuitively, the encoding lifts the persistence and data logic layers to the control layer, resulting in a “pristine” ν -CPN. To show behavioral correspondence, one should make sure that states of $\Gamma_{\mathcal{B}}^{s_0}$ and $\Gamma_{\mathcal{N}}^{M_0}$ are comparable. This can be achieved by slightly modifying the notion of weak bisimulation in such a way that, for each $\langle \langle \mathcal{I}, m \rangle, M \rangle \in wb$, we compare elements stored in $\mathcal{I}$ only with their “control counterparts” in P_{rel} of M , whereas $m \subseteq M$. Moreover, we assume that states of $\Gamma_{\mathcal{N}}^{M_0}$ are restricted only to places in $P \cup P_{rel}$, that is, each marking M shall reveal tokens stored only in P and P_{rel} , and that when constructing $\Gamma_{\mathcal{N}}^{M_0}$ all the auxiliary transitions of $\mathcal{N}$ (i.e., all the transitions within the grey lane in Figure 5) are going to be labeled with ϵ . Note that such an extended definition allows to establish equivalence not only in terms of behaviors of two systems, but also in terms of their (data) content.

4 Conclusions

We have shown that the large and relevant fragment of DB-nets employing unions of conjunctive queries with negative filters as database query language, can be faithfully encoded into a special class of Coloured Petri nets with transition priorities. Since the encoding is based on a constructive technique that can be readily implemented, the next step is to incorporate the encoding into the DB-net extension of CPN Tools [13], in turn making it possible to make the state-space construction mechanisms available in CPN Tools also applicable to DB-nets. It must be noted that, due to the presence of data ranging over infinite colour

domains, the resulting state-space is infinite in general. However, in the case of state-bounded DB-nets [11], that is, DB-nets for which each marking contains boundedly many tokens and boundedly many database tuples, a faithful abstract state space can be actually constructed using the same approach presented in [2]. Interestingly, this can be readily implemented by replacing the ML code snippet dealing with fresh value injection with a slight variant that recycles, when possible, old data values that were mentioned in a previous marking but are currently not present anymore.

References

1. Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison Wesley (1995)
2. Calvanese, D., De Giacomo, G., Montali, M., Patrizi, F.: First-order mu-calculus over generic transition systems and applications to the situation calculus. *Inf. and Comp.* (2017)
3. Calvanese, D., De Giacomo, G., Montali, M.: Foundations of data aware process analysis: A database theory perspective. In: *Proc. of PODS* (2013)
4. Drescher, C., Thielscher, M.: A fluent calculus semantics for ADL with plan constraints. *LNCS*, vol. 5293, pp. 140–152. Springer (2008)
5. Hull, R.: Artifact-centric business process models: Brief survey of research results and challenges. In: *Proc. of ODBASE*. pp. 1152–1163 (2008)
6. Jensen, K., Kristensen, L.M.: Coloured Petri Nets - Modelling and Validation of Concurrent Systems. Springer (2009)
7. de Leoni, M., Felli, P., Montali, M.: A holistic approach for soundness verification of decision-aware process models. In: *Proc. of ER*. *LNCS*, vol. 11157, pp. 219–235. Springer (2018)
8. Libkin, L.: Elements of Finite Model Theory, *LNCS*, vol. 7360, chap. Fixed Point Logics and Complexity Classes. Springer (2004)
9. Milner, R.: Communication and Concurrency. Prentice-Hall, Inc. (1989)
10. Montali, M., Rivkin, A.: Model checking petri nets with names using data-centric dynamic systems. *Formal Aspects of Computing* pp. 1–27 (2016)
11. Montali, M., Rivkin, A.: DB-Nets: on the marriage of colored Petri Nets and relational databases. *Trans. of Petri Nets and other Models of Concurrency* **28**(4) (2017)
12. Reichert, M.: Process and data: Two sides of the same coin? In: *Proc. of OTM*. pp. 2–19 (2012)
13. Ritter, D., Rinderle-Ma, S., Montali, M., Rivkin, A., Sinha, A.: Formalizing application integration patterns. pp. 11–20 (2018)
14. Rosa-Velardo, F., de Frutos-Escrig, D.: Decidability and complexity of petri nets with unordered data. *Theor. Comput. Sci.* **412**(34), 4439–4451 (2011)
15. Triebel, M., Sürmeli, J.: Homogeneous equations of algebraic petri nets. In: *Proc. of CONCUR*. pp. 1–14. *LNCS*, Springer (2016)
16. Westergaard, M., Verbeek, H.M.W.E.: Efficient implementation of prioritized transitions for high-level petri nets. In: *Proc. of PNSE*. pp. 27–41 (2011)