



UvA-DARE (Digital Academic Repository)

A Petri Net-Based Notation for Normative Modeling: Evaluation on Deontic Paradoxes

Sileno, G.; Boer, A.; van Engers, T.

DOI

[10.1007/978-3-030-00178-0_6](https://doi.org/10.1007/978-3-030-00178-0_6)

Publication date

2018

Document Version

Final published version

Published in

AI Approaches to the Complexity of Legal Systems

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Sileno, G., Boer, A., & van Engers, T. (2018). A Petri Net-Based Notation for Normative Modeling: Evaluation on Deontic Paradoxes. In U. Pagallo, M. Palmirani, P. Casanovas, G. Sartor, & S. Villata (Eds.), *AI Approaches to the Complexity of Legal Systems: AICOL International Workshops 2015-2017: AICOL-VI@JURIX 2015, AICOL-VII@EKAW 2016, AICOL-VIII@JURIX 2016, AICOL-IX@ICAIL 2017, and AICOL-X@JURIX 2017 : revised selected papers* (pp. 89-104). (Lecture Notes in Computer Science; Vol. 10791), (Lecture Notes in Artificial Intelligence). Springer. https://doi.org/10.1007/978-3-030-00178-0_6

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)



A Petri Net-Based Notation for Normative Modeling: Evaluation on Deontic Paradoxes

Giovanni Sileno^{1,2(✉)}, Alexander Boer¹, and Tom van Engers¹

¹ Leibniz Center for Law, University of Amsterdam, Amsterdam, The Netherlands
g.sileno@uva.nl

² LTCI, Télécom ParisTech, Université Paris-Saclay, Paris, France

Abstract. Developing systems operating in alignment with norms is not a straightforward endeavour. Part of the problems derive from the suggestion that law concerns a system of norms, which, in abstract, in a fixed point in time, could be approached and expressed atemporally, but, when it is contextualized and applied, it deals with a continuous flow of events modifying the normative directives as well. The paper presents an alternative approach to some of these problems, exemplified by well-known deontic puzzles, by extending the Petri net notation, most common in process modeling, to Logic Programming Petri Nets. The resulting visual formalism represents in an integrated, yet distinct fashion, procedural and declarative aspects of the system under study, including normative ones.

Keywords: Contrary-to-duty · Deontic puzzles · Logic Programming Petri nets · Normative modeling

1 Introduction

Puzzles and paradoxes are tools for bringing conceptualizations to their boundary conditions, and are therefore relevant for testing formal notations. The deontic logic community has devoted special attention to *paradoxes* constructed with *contrary-to-duty* (CTD) structures (see e.g. [3, 4, 9, 11, 17, 22]). These are “paradoxes” because, although the normative statements look plausible in natural language, when each sentence is formalized in *standard deontic logic* (SDL)—the *paper tiger* of normative modeling—either the set of formulas is inconsistent, or one of the formulas is a logical consequence of another formula (e.g. [2]).

Definition 1 (Contrary to Duty). *A contrary-to-duty (CTD) structure is a situation in which a primary obligation exists, and with its violation, a secondary obligation comes into existence.*

The importance of CTDs lies in more than just their theoretical aspects: CTDs are fundamental to normative modeling, because they are at the base

of *compensatory norms*, prototypical in e.g. contracts. The problem carries definite applicative concerns. This intuitively simple structure produces complex structures of obligations, prohibitions or permissions applying to sequences of violations or satisfactions relative to the conduct of agents in a regulated social system. Moreover, the secondary obligation may directly contradict the primary obligation. In this case, the conflict *has* to be solved to decide a course of action (i.e. “I am obliged, but I am forbidden, so what should I do?”). This consideration highlights the problem of specifying and treating preferences between idealities, or, in more agentive terms, priorities between commitments.

In this work, we focus on testing CTD structures on Logic Programming Petri Nets (LPPN).¹ This modeling notation has been introduced in [25] with the purpose of integrating and distinguishing in the same visual representation *declarative* aspects (concerning terminology, ontological constraints, normative directives, etc.) and *procedural* aspects (mechanisms, processes, courses of actions, etc.) of the reference system. The resulting common representational ground is proposed as a basis to support a continuous re-alignment in administrative organizations of representations of *law* (norms), of *implementations of law* (services as business processes), and of *action* (behavioural scripts, use cases, possibly intentionally characterized). The notation aims therefore to cover a wider class of models (business processes embedded with normative positions, scenarios issued from narratives, agent scripts, etc.) than what usually studied by deontic logic. Furthermore, beside its visual power (in principle increasing its accessibility), it enjoys computational properties as distributed computation, i.e. it does not require the reference to a global state.²

The connection of normative modeling with Petri nets is not new; see e.g. [18,19], and more recently [23]; however, these works mainly focus on events and factual conditions, overlooking normative characterizations. Other works, such as [24], have proposed using Petri nets to formalize contracts; with respect to these, the present proposal is more specific, as it focuses on minimal CTDs (through the lens of deontic puzzles), but also more general, as it considers the inclusion of declarative bindings in the model.

The paper is organized as follows. In Sect. 2, we will present *Logic Programming Petri Nets* (LPPN), delineating the notation and introducing informally a simplified version of its semantics. In Sect. 3, we will review a series of examples, all, save the first, copied from the literature. For each of them, we will propose models of the corresponding scenarios, and attempt to clarify part of the issues encountered in SDL. Discussion and further developments end the paper. A formalization of the propositional version of LPPN can be found in the appendix.

¹ Prototypes of LPPN interpreters are available on <http://github.com/s113n0/pypneu> and <http://github.com/s113n0/lppneu>.

² Cf. the recent extension to standard deontic logic by Gabbay and Straßer [6] integrating reactive constructs, an approach in many aspects dual to the present proposal.

2 Logic Programming Petri Nets

Petri nets are a simple, yet effective computational modeling representation featuring an intuitive visualization (see Fig. 1). They consist in directed, bipartite graphs with two types of nodes: *places* (visually represented with circles) and *transitions* (with boxes). A place can be connected only to transitions and vice-versa. One or more *tokens* (dots) can reside in each place. The execution of Petri nets is also named “token game”: transitions *fire* by consuming tokens from their input places and producing tokens in their output places.³




Fig. 1. Example of a Petri net and of its execution (but also of a LPPN *procedural component* when labels are propositions).

Despite their widespread use in computer science, electronics, business process modeling and biology, Petri nets are generally considered not to be enough expressive for reasoning purposes; in effect, they do not refer explicitly to any informational or representational concept. In their simplest form, tokens are indistinct, and do not transport any data. Nevertheless, usually modelers introduce labels to set up a correspondence between the *modeling* entities and the *modeled* entities. This practice enables them to read the results of a model execution in reference to the modeled system, and therefore it becomes *functional to the use* of the notation, although it is not a requirement for the execution in itself. Further interaction is possible if these labels are processed according an additional formalism, as for instance with the *Coloured Petri Net* (CPN) notation [13], which, for many aspects, is a descendant of *Predicate/Transition Nets* [7]. If its expressiveness and wide application provide reasons for its adoption, the CPN notation introduces many details which are unimportant in our setting (e.g. expressions on arcs); more importantly, it still misses the requirement of processing declarative bindings, necessary, for instance, to model terminological relationships. We opted therefore for an alternative notation.

Whereas Petri nets specify *procedural* mechanisms, LPPNs extend those (a) with Prolog-like *literals* as labels, attached on places and transitions; (b) with nodes specifying (logic) *declarative bindings* on places and on transitions.

³ For an overview on the general properties of Petri nets see e.g. [21].

The notation builds upon the intuition that places and transitions mirror the common-sense distinction between *objects* and *events* (e.g. [1]), roughly reflecting the use of *noun/verb* categories in language [14]: the procedural components can be used to model *transient* aspects of the system in focus; the declarative components to model *steady state* aspects, i.e. those on which the transient is irrelevant or does not make sense (e.g. terminology, ontological constraints, etc.).

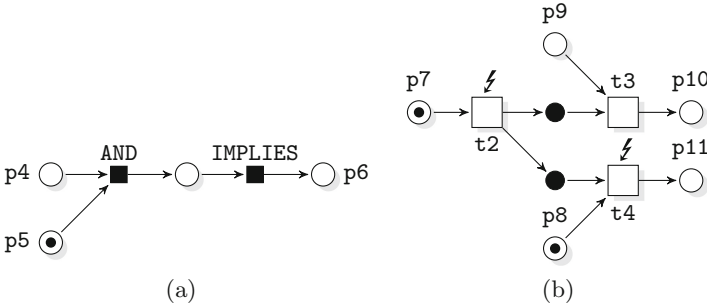


Fig. 2. Examples of LPPN *declarative components*: (a) defined on places, corresponding to the Prolog/ASP code: $p6 \text{ :- } p4, p5. p5.$ (b) defined on transitions, instantaneously propagating the firing where possible (the *IMPLIES* label on black circles is left implicit).

In this paper, for simplicity, we will consider only propositional labeling; with this assumption, the execution model of the LPPN procedural component is the same of *Condition/Event* nets, i.e. Petri nets whose places are not allowed to contain more than one token. For this reason, the Petri net in Fig. 1 can be interpreted as an example of LPPN specifying a procedural mechanism. However, the LPPN notation introduces also logic operator nodes (or *l-nodes*), which apply on places or on transitions. An example of a sub-net with *l-nodes* for places (small black squares) is given in Fig. 2a. These are used to create logic compositions of places (via operators as *NEG*, *AND*, *OR*, etc.), or to specify logic interdependencies (via the logic conditional *IMPLIES*). Similarly, transitions may be connected declaratively via *l-nodes* for transitions (black circles) as in Fig. 2b. These connections may be interpreted as channels enabling *instantaneous propagation* of firing. In this case, it is not relevant to introduce operators as *AND* because, for the interleaving semantics, only one source transition may fire per step. To simplify the visual burden, we might leave the *IMPLIES* label implicit, exploiting the sense of the arrow to specify the direction of the relation. Operationally, these declarative components are treated integrating the *stable model semantics* used in *answer set programming* (ASP) [15]. This was a natural choice because process execution exhibits a prototypical ‘forward’ nature, and ASP can be interpreted as providing forward chaining. A formalization of propositional LPPNs can be found in the Appendix.

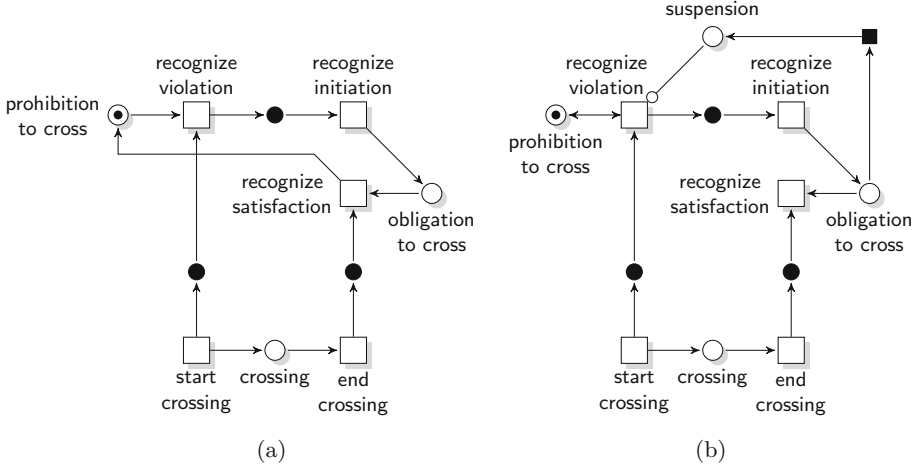


Fig. 3. A minimal, conflicting and event-based *contrary-to-duty* (CTD) structure: a secondary obligation is created after the violation of a primary obligation, and in conflict with it. The CTD may be interpreted: (a) as an *exception*; or (b) as *overriding* the primary obligation. The second is the most accepted option.

3 Deontic Exercises

3.1 Crossing or Not Crossing?

Let us start from this minimal, conflicting CTD structure:

You are forbidden to cross the road.

If you are crossing the road, (you have to) cross the road!

This rule of conduct is perfectly plausible: most parents say something similar to their children at some moment. However, its translation in basic deontic logic is not direct. The text suggests, in effect, an underlying model in terms of action: a state-based interpretation would miss the implicit initiation/termination events that make the action-wise prescription sound, and namely:

You are forbidden to cross the road.

If you have started to cross the road, you are obliged to finish crossing.

This *transitional* aspect can be easily mapped on a LPPN, separating the experiential world from the institutional world, with the second synchronized to the first via *constituting links* determining what *counts as* violation or satisfaction.⁴

⁴ With respect to constitutive rules, the LPPN notation enables to easily distinguish *classificatory constitutive rules* (e.g. “a bike counts as a vehicle”) from *constitutive event rules* (e.g. “raising a hand counts as making a bid”), as they are modeled respectively using black boxes or black circles. Most formalizations of constitutive rules consider only on the first aspect (e.g. [10]), cf. the overview in [27].

In principle, two modeling options are available in regard to the secondary obligation; it can interpreted:

- as an *exception*, thus temporarily retracting the primary obligation (Fig. 3a);
- as *overriding* the primary obligation, which persists concurrently (Fig. 3b).

The second option requires an additional treatment, because it brings two contrary/opposite positions to hold concurrently. Similarly to what suggested in the literature, this can be solved introducing an explicit *ordering* between positions, which depends on how close to *ideal* is the world/context they are referring to (see e.g. [16, 22]). In the proposed Petri net an aspect of this mechanism—the fact that the secondary obligation is put in force *in response to the violation* of the primary one—is already reified in the topology. To capture the remaining part, i.e. that the second is *contextually overriding* the first, we need to order them in the opposite sense: the last obligation created is the one with most priority and should be the only active, *suspending* the previous ones. This can be done introducing an *inhibiting arc* (visualized in Fig. 3b as an arrow with a circle-shaped head).⁵ The resulting design can be seen as a model of *salience*.




Fig. 4. Gentle murderer case.

⁵ Inhibiting arcs goes from places to transitions. If the input place of an inhibiting arc is occupied, its output transition is disabled.

3.2 Gentle Murderer

The previous CTD model gives us the basic instruments to proceed. Let us start from the classic case of the “gentle murderer”, given by Forrester [5]:

*It is forbidden to kill,
but if one kills, one ought to kill gently.*

This example is very similar to the previous one, except that the target of the secondary obligation is subsumed by the target of the first one. Because our notation explicitly accounts for a declarative dimension for events, we can directly map this relation (Fig. 4).




Fig. 5. White fence case.

3.3 White Fence

Now, we consider a static and extended variation proposed by Prakken and Sergot [22], the “white fence” case:

*There must be no fence.
If there is a fence, it must be a white fence.
If the cottage is by the sea, there may be a fence.*

This example shows the importance of distinguishing exceptions from overriding effects due to CTDs (and therefore supports the second interpretation of CTD given in Sect. 3.1). In principle, a rule specifies a CTD if its premise is the negation of the target of the obligation in the consequent of another rule. A rule specifies an exception if it has as consequent the negation of the consequent of another rule, and it has a lower *priority* than the first one (exceptions are by definition subordinate to some *normal* conditions). In effect, the two rules can

be read as referring to a *priority-based representation* [26]. Considering part of the “white fence” case in propositional form, we have:

$$Forb(fence) \quad sea \rightarrow Perm(fence)$$

which can be translated to the corresponding *constraint-based* representation:

$$\neg sea \rightarrow Forb(fence) \quad sea \rightarrow Perm(fence)$$

This treatment gives a hint as to how to deal with *exceptions*—that is, it helps make explicit an enchaining of negations of the premises following the inverse ordering of salience. The fastest solution to avoiding conflicts in the case of belief revision is to not reify directly the *default* position (in this case, prohibition against having a fence), but to generate it through a *default rule* [26]:

$$not\ sea \rightarrow Forb(fence)$$

The resulting model is illustrated in Fig. 5.

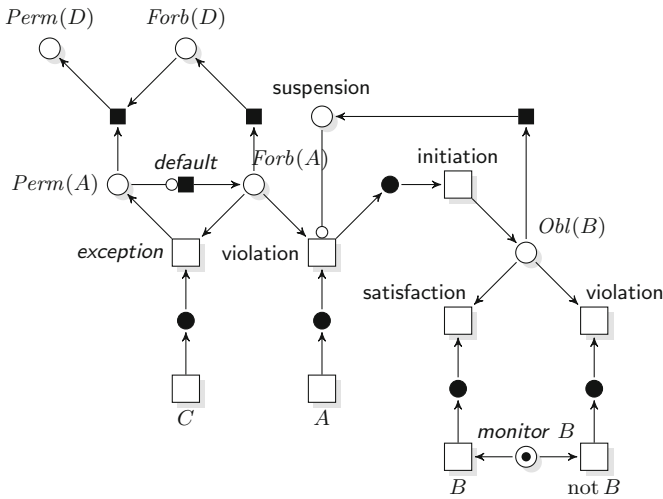


Fig. 6. Privacy act case.

3.4 Privacy Act

Recently, Governatori [8] has proposed the case of a Privacy Act (fictional, but based on actual Australian normative provisions):

- i. *The collection of personal information is forbidden, unless acting on a court order authorising it.*

- ii. *The destruction of illegally collected personal information before accessing it is a defence against the illegal collection of the personal information.*
- iii. *The collection of medical information is forbidden, unless the entity collecting the medical information is permitted to collect personal information.*

The following deontic interpretation is proposed:

- i. Forbidden A. If C, then Permitted A.
- ii. If Forbidden A and A, then Obligatory B.
- iii. Forbidden D. If Permitted A, then Permitted D.

A, B, C, and D in this specific case are actions; (ii) specifies a CTD, (i) and (iii) provide rules based on a priority-based representation. As before, we extract explicitly the defaults. The negation of permission of A in (iii) can be interpreted as the prohibition of A, thus converging to the default in (i). For completeness, we have reported in Fig. 6 the monitoring place from which the event B or not B is recognized.

3.5 Detachment Principles

In the deontic logic literature, two types of “detachment principles” are recognized as relevant. The first is called *factual detachment* (FD):

$$p \wedge Obl(q|p) \rightarrow Obl(q) \quad (1)$$

The second is known as *deontic detachment* (DD):

$$Obl(p) \wedge Obl(q|p) \rightarrow Obl(q) \quad (2)$$

In our framework, a conditional directive or commitment is seen as a *susceptibility* to a condition that *creates* or *implies* the directive depending upon whether the connective is a *causal* or *logical dependence*. The two principles can be then translated using the LPPN notation. Focusing on the logical dependence case, the result is seen in Fig. 7. The pictures show that the first principle is satisfied by the notation semantics; on the contrary, the second principle, which is based on an *anticipation* of the normal conditional, is not satisfied.

3.6 Derived Obligation

Consider these two sentences:

- *Bob’s promise to meet you commits him to meeting you.*
- *It is obligatory that if Bob promises to meet you, he does so.*

Although in natural language their difference is arguable, in the literature they have been formalized using two distinct deontic formulations:

- $p \rightarrow Obl(m)$
- $Obl(p \rightarrow m)$




Fig. 7. Factual and deontic detachments.

What is the difference between the two formulations in our framework? The second formula can be translated with the consideration that the obligation of *something* consists of two recognition rules about satisfaction and violation, by default anchored respectively to *something* and to \neg *something*. As an object, the conditional within the obligation can be transformed using the *material implication*.⁶ The result are reported in Fig. 8. As we see in the picture, both models are violated in the same situation (p and $\neg m$); however, the second includes the recognition of a satisfied situation not accounted for in the first ($\neg p$). In other words, the first derived obligation precisely discriminates the elements producing the violation. The second takes an explicit position also on the satisfying elements.

3.7 Chisholm’s Paradox

At this point, we can finally model the “paradox” proposed by Chisholm [4]:

It ought to be that Jones goes (to the assistance of his neighbors).
It ought to be that if Jones goes, then he tells them he is coming.
If Jones doesn’t go, then he ought not tell them he is coming.
Jones doesn’t go.

This was seen as a paradox, because if we model it as:

- i. $Obl(go)$
- ii. $Obl(go \rightarrow tell)$
- iii. $\neg go \rightarrow Forb(tell)$
- iv. $\neg go$

⁶ The specific example from which we started is not based on a logic conditional, but on a causal connective, at least in the case of “if Bob promises to meet you, then he does so”. In this case, the use of material implication is not a perfect fit, as the temporal shift between the promise and the meeting falsifies the derived constraint, at least on a transient basis. On a steady state analysis, however, this simplification may be applied.

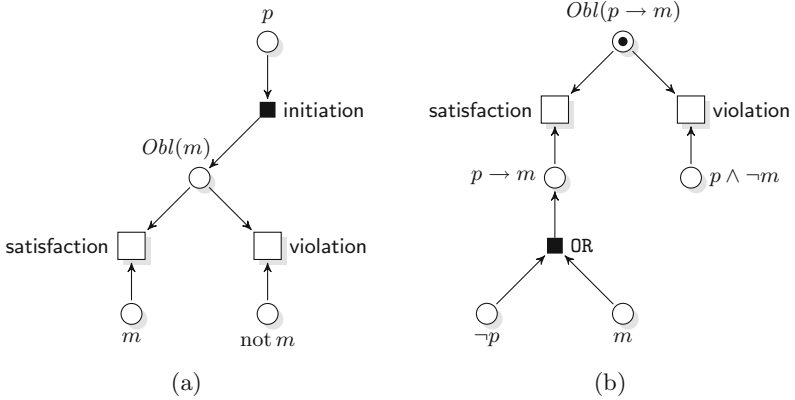


Fig. 8. Derived obligation modeled as $p \rightarrow Obl(m)$ or as $Obl(p \rightarrow m)$.

and we apply both deontic and factual detachments, we find an inconsistency. More precisely, from (i) and (ii), using of deontic detachment, we derive $Obl(tell)$, while from (iii) and (iv), using factual detachment, we derive $Obl(\neg tell)$. However, representing this model using our notation as in Fig. 9, we do not find any specific issue. The fact $\neg go$ satisfies $Obl(go \rightarrow tell)$, but violates $Obl(go)$, and for this reason, $Forb(tell)$ is instantiated. Depending on whether $tell$ becomes true, this may or may not be violated.

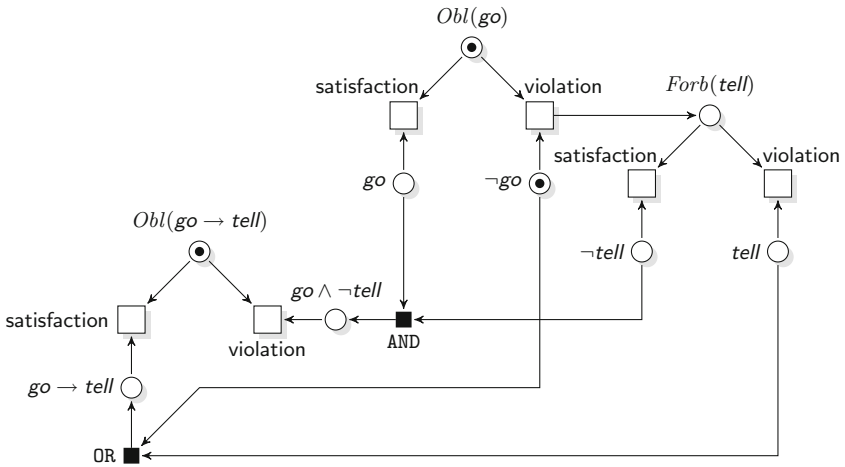


Fig. 9. Chisholm's paradox case.




Fig. 10. Residential neighbourhood case.

3.8 Residential Neighbourhood

As we observed above, the difference in natural language between the two types of derived obligations is arguable. In order to show what would happen if we interpret the two as conditional obligations, we consider the non-agent version of the Chisholm’s paradox proposed by Prakken and Sergot [22]:

There must be no dog.
If there is no dog, there must be no warning sign.
If there is a dog, there must be a warning sign.
There is a dog.

Written in this way, the modeling is straightforward (Fig. 10).

4 Discussion and Further Developments

At superficial level, the paper presented examples of application of a notation introduced for wider modeling purposes [25]. This served as an important exercise to (partially) evaluate its practical functionality, as one may capture the subtle problems related to modeling in a certain domain (in this case normative modeling) only by approaching the specific issues that have been raised by practitioners and scholars in that domain. The introduction of *defaults*, *exceptions*, and *suspensions* in the notation presented here is a preliminary proposal. Some of these aspects have been treated more in more detail elsewhere (e.g. *defaults* in priority-based rule-bases [26], proceeding along [12]), while others require further investigation (e.g. *suspension*, cf. [20]).

Nevertheless, the exercise yields concrete practical and theoretical results. From a practical point of view, it makes a case supporting normative modeling with notations similar to those used in business process modeling, thus potentially facilitating cross-fertilization between theoretical to operational settings. From a theoretical point of view, we observed for instance that our conceptual framework does not entail the *deontic detachment* principle, hinting at a more general *minimal commitment* taken by the notation. In other words, i.e. the notation does not provide any rule *a priori* to conclude whether $Obl(A) \wedge Obl(B)$ is the same as $Obl(A \wedge B)$. This neutral starting point can be used to evaluate the

alternative impact of different axioms proposed in the literature, in affinity with approaches like *input/output logic* [17].

More importantly, while working on these exercises, we appreciated the crucial interplay between static and dynamic aspects (one of the issues underlying many deontic puzzles). The LPPN notation, requiring the explicitation of procedural and declarative aspects, highly facilitated this task, but our exercise suggests further research on the modeling methodology. For instance, in Chisholm’s paradox, we considered *go* and *tell* as labels of places, but strictly speaking, they should be attached to transitions (as in the *privacy act* case). With this choice, we would require making explicit the *occurrence* of the events as places, in order to evaluate the material implication. Does this simplification hint at a more general pattern? We modeled the nodes concerning violation and satisfaction as l-nodes only when their transformational nature was certain. In general, however, they may be *transformational* (when they simply identify whether a violation or satisfaction holds at the moment) or *reactive* (when they reify the fact that a violation occurred in that moment, or when they cause any change on the inputs, e.g. removing the obligation)—note again the interplay between static and dynamic aspects. In the future, an investigation of boundary cases may help in formulating a more general theory on how to decide upon the level of abstraction, and about whether it is possible to identify or elaborate general patterns, that, depending on circumstances, are read in one form or in the other.

A Formalization

Here we present a simplified version of the LPPN notation considering only a *propositional* labeling. We start from the definition of propositional literals derived from ASP [15], accounting for strong and default negation.

Definition 2 (Literal and Extended literals). *Given a set of propositional atoms A , the set of literals $L = L^+ \cup L^-$ consists of positive literals (atoms) $L^+ = A$, negative literals (negated atoms) $L^- = \{-a \mid a \in A\}$, where ‘ $-$ ’ stands for strong negation.⁷ The set of extended literals $L^* = L \cup L^{not}$ consists of literals and default negation literals $L^{not} = \{notl \mid l \in L\}$, where ‘not’ stands for default negation.⁸*

We denote the basic topology of a Petri net as a procedural net.

Definition 3 (Procedural net). *A procedural net is a bipartite directed graph connecting two finite sets of nodes, called places and transitions. It can be written as $N = \langle P, T, E \rangle$, where $P = \{p_1, \dots, p_n\}$ is the set of place nodes; $T = \{t_1, \dots, t_m\}$ is the set of transition nodes; $E = E^+ \cup E^-$ is the set of arcs connecting them: E^+ from transitions to places, E^- from places to transitions.*

⁷ Strong negation is used to reify an explicitly false situation (e.g. “It does not rain”).

⁸ Default negation is used to reify a situation in which something cannot be retrieved/inferred (e.g. ‘It is unknown whether it rains or not’).

LPPNs consists of three components: a procedural net specifying causal or temporal relationships, and two declarative nets specifying respectively logical dependencies at the level of objects or ongoing events (on places), and on impulse events (on transitions). Furthermore, *propositional* LPPNs build upon a boolean marking on places (like *condition/event* nets).

Definition 4 (Propositional Logic Programming Petri Net). A propositional Logic Programming Petri Net $LPPN_{\text{prop}}$ is a Petri Net whose places and transitions are labeled with literals, enriched with declarative nets of places and of transitions. It is defined by the following components:

- $\langle P, T, PE \rangle$ is a procedural net; PE stands for procedural edges;
- $C_P : P \rightarrow L^*$ and $C_T : T \rightarrow L$ are labeling functions, associating literals respectively to places and to transitions;
- $OP = \{\neg, -, \wedge, \vee, \rightarrow, \leftrightarrow, \dots\}$ is a set of logic operators.
- LP and LT are sets of logic operator nodes (in the following called l-nodes) respectively for places and for transitions.
- $C_{LP} : LP \rightarrow OP$ maps each l-node for places to a logic operator; similarly, $C_{LT} : LT \rightarrow OP$ does the same for l-nodes for transitions.
- $DE_{LP} = DE_{LP}^+ \cup DE_{LP}^-$ is the set of arcs connecting l-nodes for places to places; similarly, $DE_{LT} = DE_{LT}^+ \cup DE_{LT}^-$ for l-nodes for transitions and transitions.⁹
- $M : P \rightarrow \{0, 1\}$ returns the marking of a place, i.e. whether the place contains (1) or does not contain (0) a token.

Note that if $LP \cup LT = \emptyset$, we have a *strictly procedural* $LPPN_{\text{prop}}$, i.e. a standard binary Petri net. If $T = \emptyset$, we have a *strictly declarative* $LPPN_{\text{prop}}$, that can be directly mapped to an ASP program.

With respect to the *operational semantics*, the execution cycle of a LPPN consists of four steps: (1) given a “source” marking M , the bindings of the declarative net of places entail a “ground” marking M^* ; (2) an enabled transition is selected to *pre-fire*, so determining a “source” *transition-event* e ; (3) the bindings of the declarative net of transitions entail all propagations of this event, obtaining a set of *transition-events*, also denoted as the “ground” *event-marking* E^* ; (4) all transition-events are fired, producing and consuming the relative tokens. The steps (1) and (3) are processed by an ASP solver: the declarative net of places (respectively transitions) is translated as *rules*, tokens (transition-events) are reified as *facts*; the ASP solver takes as input the resulting program and, if satisfiable, it provides as output one or more ground marking (one or more sets transition-events to be fired). For the steps (2) and (4), the operational semantics distinguishes the *external* firings (started by the execution) from the *internal* firing, immediately propagated (triggered by the declarative net of transitions).

⁹ Note that $DE_{LT}^- \subseteq (T \cup P) \times LT$, i.e. these edges go from transitions *and* places (modeling contextual conditions) to l-nodes for transitions.

Definition 5 (Enabled transition). A transition t is enabled in a ground marking M^* if a token is available for each input places:

$$Enabled(t) \equiv \forall p_i \in \bullet t, M^*(p) = 1$$

Similarly to what marking is for places, we consider an *event-marking* for transitions $E : T \rightarrow \{0, 1\}$. $E(t) = 1$ if the transition t produces a transition-event e . Each step s has a “source” event-marking E .

Definition 6 (Pre-firing). An enabled transition t pre-fires at a step s if it selected to produce a transition-event:

$$\forall t \in Enabled(T) : t \text{ pre-fires} \equiv E(t) = 1$$

As we apply an *interleaving semantics* for the pre-firing, the interpreter selects only one transition to pre-fire per step; for any other t' , $E(t') = 0$.

Definition 7 (Firing). An enabled transition t fires by propagation, consuming a token from each input place, and forging a token in each output place:

$$\forall t \in Enabled(T) : t \text{ fires} \equiv \\ E^*(t) = 1 \leftrightarrow \forall p_i \in \bullet t : M'(p_i) = 0 \wedge \forall p_o \in t \bullet : M'(p_o) = 1$$

References

1. Breuker, J., Hoekstra, R.: Core concepts of law: taking common-sense seriously. In: Proceedings of Formal Ontologies in Information (2004)
2. Broersen, J., van der Torre, L.: Ten problems of deontic logic and normative reasoning in computer science. In: Bezhanishvili, N., Goranko, V. (eds.) ESSLLI 2010-2011. LNCS, vol. 7388, pp. 55–88. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31485-8_2
3. Carmo, J., Jones, A.: Deontic logic and contrary-to-duties. In: Gabbay, D.M., Guenther, F. (eds.) Handbook of Philosophical Logic, vol. 8, pp. 265–343. Springer, Dordrecht (2002). https://doi.org/10.1007/978-94-010-0387-2_4
4. Chisholm, R.M.: Contrary-to-duty imperatives and deontic logic. *Analysis* **24**(2), 33–36 (1963)
5. Forrester, J.W.: Gentle murder, or the adverbial Samaritan. *J. Philos.* **81**(4), 193–197 (1984)
6. Gabbay, D.M., Straßer, C.: Reactive standard deontic logic. *J. Log. Comput.* **25**(1), 117–157 (2015)
7. Genrich, H.J.: Predicate/transition nets. In: Brauer, W., Reisig, W., Rozenberg, G. (eds.) ACPN 1986. LNCS, vol. 254, pp. 207–247. Springer, Heidelberg (1987). https://doi.org/10.1007/978-3-540-47919-2_9
8. Governatori, G.: Thou shalt is not you will. Technical report, NICTA (2015)
9. Governatori, G., Rotolo, A.: Logic of violations: a gentzen system for reasoning with contrary-to-duty obligations. *Australas. J. Log.* **4**, 193–215 (2006)
10. Grossi, D., Meyer, J.J.C., Dignum, F.: Classificatory aspects of counts-as: an analysis in modal logic. *J. Log. Comput.* **16**(5), 613–643 (2006)

11. Hansen, J., Pigozzi, G., Van Der Torre, L.: Ten philosophical problems in deontic logic. *Normative Multi-agent*, pp. 1–26 (2007)
12. Horty, J.F.: Rules and reasons in the theory of precedent. *Legal Theory* **17**(1), 1–33 (2011)
13. Jensen, K.: *Coloured Petri Nets: Basic Concepts, Analysis Methods and Practical Use*. Springer, Heidelberg (1996). <https://doi.org/10.1007/978-3-662-03241-1>
14. Kemmerer, D., Eggleston, A.: Nouns and verbs in the brain: implications of linguistic typology for cognitive neuroscience. *Lingua* **120**(12), 2686–2690 (2010)
15. Lifschitz, V.: What is answer set programming? In: *Proceedings of the 22th AAAI Conference on Artificial Intelligence* (2008)
16. Makinson, D.: Five faces of minimality. *Stud. Log.* **52**, 339–379 (1993)
17. Makinson, D., Van Der Torre, L.: Input/output logics. *J. Philos. Log.* **29**, 383–408 (2000)
18. Meldman, J., Fox, S.: Concise petri nets and their use in modeling the social work (Scotland) Act 1968. *Emory Law J.* **30**, 583–630 (1981)
19. Meldman, J., Holt, A.: Petri nets and legal systems. *Jurimetr. J.* **12**(2), 65–75 (1971)
20. Meneguzzi, F., Telang, P., Singh, M.: A first-order formalization of commitments and goals for planning. In: *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pp. 697–703 (2013)
21. Murata, T.: Petri nets: properties, analysis and applications. *Proc. IEEE* **77**(4), 541–580 (1989)
22. Prakken, H., Sergot, M.: Contrary-to-duty obligations. *Stud. Log.* **57**(1), 91–115 (1996)
23. Purvis, M.A.: *Dynamic modelling of legal processes with petri nets*. Ph.D. thesis, University of Otago (1998)
24. Raskin, J.F., Tan, Y.H., van der Torre, L.: How to model normative behavior in Petri nets. In: *Proceedings of the 2nd ModelAge: Workshop on Formal Models of Agents*, pp. 223–241 (1996)
25. Sileno, G.: *Aligning law and action*. Ph.D. thesis, University of Amsterdam (2016)
26. Sileno, G., Boer, A., van Engers, T.: A constructivist approach to rule bases. In: *Proceeding of the 7th International Conference on Agents and Artificial Intelligence (ICAART 2015)* (2015)
27. Sileno, G., Boer, A., van Engers, T.: Revisiting constitutive rules. In: *Proceedings of the 6th Workshop on Artificial Intelligence and the Complexity of Legal Systems (AICOL 2015)* (2015)