Quantitative Deadlock Analysis in Petri Nets using Inconsistency Measures

Elina Unruh
Department of Computer Science
University of Koblenz-Landau
Koblenz, Germany
elina@uni-koblenz.de

Patrick Delfmann
Department of Computer Science
University of Koblenz-Landau
Koblenz, Germany
delfmann@uni-koblenz.de

Matthias Thimm
Department of Computer Science
University of Koblenz-Landau
Koblenz, Germany
thimm@uni-koblenz.de

Abstract*—Petri Nets are often used to describe, execute, analyze and improve business processes. A special area of interest is the detection of possible deadlocks. Deadlocks can harm the proper execution of business processes which may lead to errors or even impossible business process execution, and, in turn, economic loss. In most cases, it is only determined whether a deadlock can occur, in order to eliminate or avoid it. For this, it is necessary to manipulate the behavior of the net, and thus change it. A so far little considered question is how the severity of the potential to encounter a deadlock can be formally investigated and assessed, without altering the net. This could be useful, for example, to assess and compare systems and processes. In a naive approach, the exhaustive calculation of all possible states of the net would be necessary to check in how many of them a deadlock occurs. This is not reasonable in most cases due to the often high to potentially indefinite number of reachable states. In this paper, approaches are developed to approximate the termination potential of an ordinary marked Petri Net through structural analysis. We do this by leveraging approaches from inconsistency measurement, which is a field within Artificial Intelligence to quantitatively assess the severity of inconsistency in formal knowledge representation formalisms. We develop six different measures for the Petri Net setting and investigate their formal properties, in particular wrt. some rationality postulates that were also adapted from the field of inconsistency measurement.

***Index Terms*—Petri nets, deadlocks, inconsistency measurement**

I. Introduction

Petri Nets are a modeling tool for the description and analysis of concurrent processes and complex distributed systems. In contrast to purely sequential models, the dependence and independence of processes can be represented explicitly and the limitation of resources can be described. Petri Nets can therefore be used in various contexts. For instance, business processes, organizational structures, operating systems, communication protocols and production control systems can be modeled appropriately. Amongst others, Petri Nets are of special interest in the field of business process automation (For an exhaustive overview on Petri Nets see [1]–[5]). Rather than just describing a business process diagrammatically, Petri Nets come with a fully formalized, built-in execution concept that makes it possible to not provide a model of a business process but also enact business processes based on that model. However, providing a formalized execution concept does not protect Petri Nets from execution problems by itself. One such execution problem is called a deadlock. Deadlocks can harm the proper execution of business processes which may lead to errors or even impossible business process execution, and, in turn, economic loss. Therefore, concepts are needed that detect deadlocks and rate their severity (see [6] and [7] for example).

In Petri Nets, a deadlock state is called *dead marking*. To the best of our knowledge, previous works on the analysis of possible dead markings focus on *existence questions* and a *qualitative analysis*. If the possibility for dead markings is detected, all effort is put into solving the problem and no quantitative evaluation of the severity of the problem is performed. However, a *quantitative* analysis of the severity of such issues can provide deeper insights about the reliability of the net. This information can be valuable, for example, for comparing or assessing different nets. Regarding the fact that the computation of all reachable states of a net is not feasible, due to the often high to potentially indefinite number of reachable states [3], [4], [8], and that from a purely structural analysis it is often not even clear whether the dead marking(s) found can actually be reached, it seems even more useful to know how much potential for dead markings there is. This information can, for example, help to decide, whether a process should be improved or rather replaced entirely.

The occurrence of dead markings is closely linked to structures in Petri Nets called *deadlocks*[1]. This connection has already been investigated and is used for the detection and avoidance of dead markings [7]. In this work, we develop measures to quantify the presence and severity of deadlocks. We approach this by using insights from the field of inconsistency measurement [9], which is concerned with assessing the severity of conflicts in logical structures. More concretely, we adapt and extend existing inconsistency measures to the setting of Petri nets. Compared to the setting of measuring inconsistency in (classical) logics, Petri nets bring about a new challenge due to their inherent *non-monotonic* behaviour: deadlocks may be introduced and/or resolved when a net is extended. In this work, we focus on measures that analyse minimal inconsistent subsets and define corresponding approaches that analyse minimal deadlocks and address this non-monotonicity aspect as well. We pursue a principle-based approach to evaluate the quality of these measures by adapting and motivating a series of

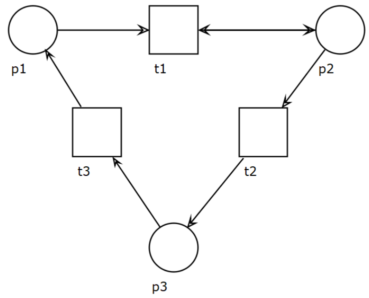

Figure 1. Petri Net

rationality postulates, i. e., desirable criteria for a meaningful account to such quantitative approaches. We finally analyse our measures wrt. their compliance to these postulates.

This work is structured as follows: At first, we explain the necessary background, i. e., we introduce Petri Nets and their properties and the field of inconsistency measurement (Section II). We then adapt some existing inconsistency measures to Petri Nets and define some rationality postulates, that are potentially desirable properties for our measures (Section IV). Finally, we discuss the relation to other works, conclude and give an outlook on what could be done in the future (Section III).

II. Background

This section gives an overview of the basic principles of Petri Nets (in particular on deadlocks and traps) and inconsistency measurement.

A. Petri Nets

Petri Nets [1]–[5] are bipartite directed graphs. The set of nodes consists of so-called places and transitions. Places are represented by circles, transitions by rectangles (or bars). A simple Petri Net is shown in Figure 1. The formal notation we will use here is inspired by the notations used in [3], [10].

Definition 1. *A Petri Net is a tuple $N = (P, T, I, O)$ with*

$$P \cap T = \emptyset \qquad\qquad I, O \subseteq P \times T.$$

*where $P = \{p_1, ..., p_n\}$ is the set of places, $T = \{t_1, ..., t_m\}$ is the set of transitions, $I = \{(p,t) \mid t \to p \text{ edge of } N\}$ is the set of edges from transitions to places (*input*), and $O = \{(p,t) \mid p \to t \text{ edge of } N\}$ is the set of edges from places to transitions (*output*).*

Definition 2. *The set ${}^\bullet p = \{t \in T \mid (p,t) \in I\}$ of transitions connected to a place p by input edges is called* pre-set *of p (see Figure 2). The set $p^\bullet = \{t \in T \mid (p,t) \in O\}$ of transitions connected to a place p by output edges is called* post-set *of p (see Figure 3). Analogous, this can be defined for transitions as well, with ${}^\bullet t = \{p \in P \mid (p,t) \in O\}$ and $t^\bullet = \{p \in P \mid (p,t) \in I\}$ and for sets of places/transitions $X \subseteq P \vee X \subseteq T$, $x \in X$ with ${}^\bullet X = \bigcup {}^\bullet x$ and $X^\bullet = \bigcup x^\bullet$ (see Figure 4 and 5).*

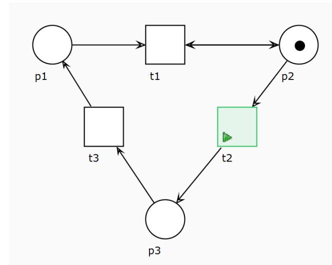

Figure 6. $t2$ is enabled

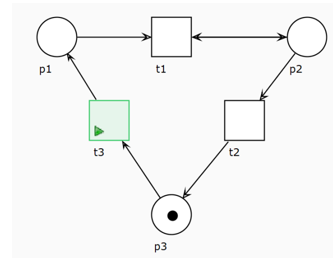

Figure 7. after $t2$ was fired

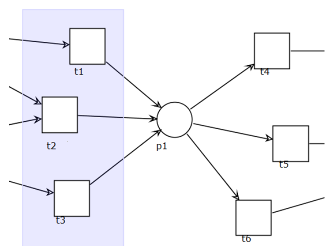

Figure 2. Pre-set of $p1$

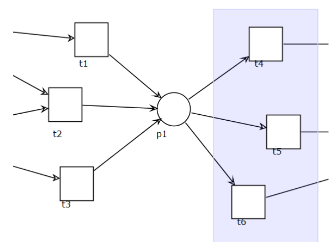

Figure 3. Post-set of $p1$

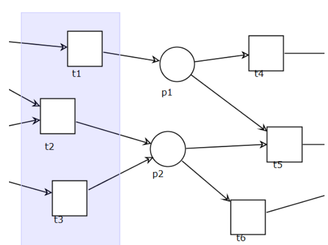

Figure 4. Pre-set of the set $\{p1, p2\}$

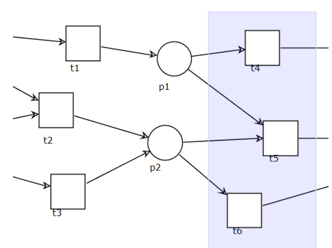

Figure 5. Post-set of the set $\{p1, p2\}$

Places can contain *tokens*, represented by dots. For ordinary marked nets [8]—the type of Petri Net we consider in this work—there is no limit to the number of tokens per place. Transitions behave like switches. If tokens are present at all places in the pre-set of a transition, the transition is called *enabled*. When a transition is fired (switched), a token is taken from each place in the pre-set and a token is added to each place in the post-set (see Figure 6 and 7).

The information about how many tokens are located at which places is called *marking* and denoted by M. Formally, M is a function $M : P \to \mathbb{N}$ where the number of tokens contained in a place p at M is referred to via $M(p)$. A marked net is represented by the tuple (N, M) and the *initial* marking is denoted by M_0.

Some further important notions are as follows:

[1]Note that from now on in this work the term deadlock is only used to refer to the homonymous structure in Petri Nets, and not as a synonym for dead marking – unless explicitly stated otherwise. The existence of deadlocks does not necessarily imply that a dead marking will or even can be reached. In the literature deadlocks are sometimes named siphons or co-traps instead.

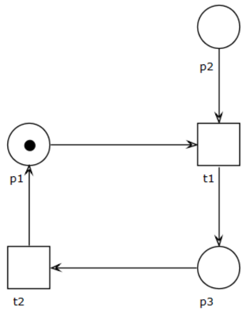

Figure 8. Dead marking

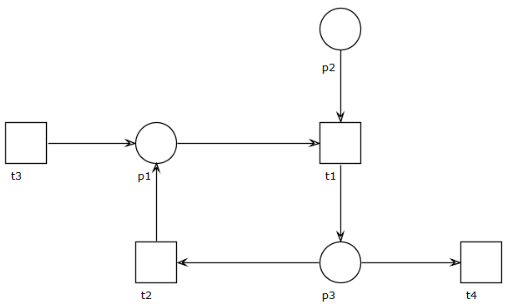

Figure 9. Deadlock-free Petri Net with source and sink transition

- A transition without input places is called *source transition* – it is always enabled and generates tokens without consuming any.
- A transition without output places is called *sink transition*. It only consumes tokens.
- A *reachable marking* is any marking that can be reached from M_0 by a finite sequence of switching transitions.
- A marking is called *dead marking* iff no transition is enabled at this marking. The net is then called *blocked*. A net that reaches a dead marking *terminates*.
- A net is *deadlock-free* iff no dead marking is reachable, i.e. every reachable marking enables at least one transition.
- A net is *live* if for every transition t and every reachable marking M, there is a reachable marking M' that activates t, i.e. every transition can be activated infinitely often.

In Figure 8 the depicted net is blocked, because no transition is enabled. In Figure 9, $t3$ is a source transition and $t4$ is a sink transition. Because $t3$ is always enabled, the net is deadlock-free. The net from Figure 10 is live, because every transition can be enabled infinitely often. In Figure 9 in contrast, transition $t1$ cannot be enabled, thus the net is not live.

B. Deadlocks and Traps

A deadlock is a subset of places such that its pre-set is completely contained in its post-set, i. e., every input transition of any place in the set is as well an output transition for some place of the set. As a consequence, if there are no tokens in the set, this cannot change, since no input transition can be enabled.

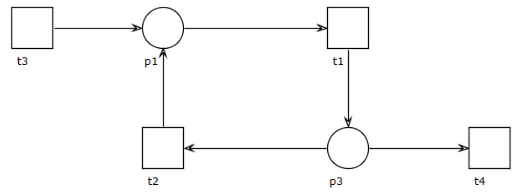

Figure 10. Live Petri Net

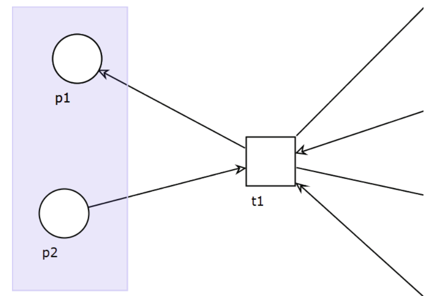

Figure 11. Part of a Petri Net containing a deadlock

Definition 3. *A set of places $P' \subseteq P$ is a* deadlock*, iff ${}^\bullet P' \subseteq P'^\bullet$.*

The set of places $\{p1, p2\}$ from Figure 11 constitutes a deadlock, since the only input transition is an output transition for one of the places as well. Other places that are connected to the same transitions have no influence on this.

Definition 4. *A set of places $P' \subset P$ is a* minimal deadlock*, iff it is a deadlock and $P'' \subset P' \implies {}^\bullet P'' \not\subseteq P''^\bullet$. The set of all minimal deadlocks of a net N is denoted by $\mathbb{D}(N)$.*

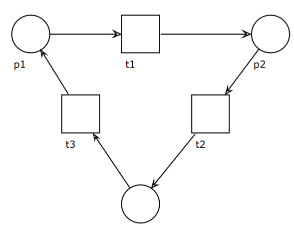

Figure 12. Minimal deadlock

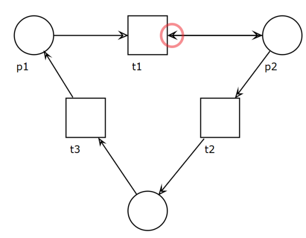

Figure 13. Non-minimal deadlock

In Figure 12 the places $p1, p2$ and $p3$ are constituting a minimal deadlock, since removing one or two places from the set would result in an input transition that is not an output transition for one of the remaining places. In Figure 13 the set $\{p1, p2, p3\}$ is a non-minimal deadlock, because $\{p2, p3\}$ is still a deadlock, as well as $\{p2\}$ alone.

A *trap* is a subset of places such that its post-set is completely contained in its pre-set, i.e. every output transition of any place in the set is as well an input transition for some place of the set.

Definition 5. *A set of places $P' \subset P$ is a* trap *iff $P'^\bullet \subseteq {}^\bullet P'$.*

Thus every transition that consumes tokens from the set generates at least one token inside the trap. Therefore, once marked, the trap remains marked (i. e. contains at least one token in total) for all reachable markings. For this reason, a deadlock containing at least one initially marked trap cannot get unmarked. Furthermore, if every deadlock in the net contains an initially marked trap, no dead marking is reachable [4]: *Consider a dead marking M_d, i.e. there are no transitions enabled in the whole net. Consider the set P_u of all unmarked places of the net. The pre-set of P_u consists of not enabled transitions, i.e. the pre-set of every transition $t \in {}^\bullet P_u$ contains at least one unmarked place (otherwise the transition would be enabled), which means all transitions from the pre-set of P_u are in the post-set as well. Thus P_u is an empty deadlock. This means, without the existence of an unmarked deadlock, the marking cannot be dead. Therefore a marked net is deadlock-free if it does not contain deadlocks without marked trap.* To verify this condition it is sufficient to check the minimal deadlocks of the net, because every non-minimal deadlock contains a minimal deadlock by definition.

In Figure 12, $\{p1, p2, p3\}$ constitutes a minimal deadlock and a trap at the same time – the deadlock contains a trap. As we can easily see, if there were tokens in one or more of the places, the deadlock/trap could not get empty, but at the same time it cannot get marked at any reachable marking when it is not marked initially. In Figure 13 on the other hand, $\{p1, p2, p3\}$ is as well deadlock and trap at once, but $\{p2, p3\}$ and $\{p2\}$ that are deadlocks as well do not contain a trap.

For more detailed information about deadlocks and traps see [4], [5], [7], [11].

C. Inconsistency Measurement

The aim of this paper is to quantitatively assess the severity imposed by deadlocks as an indicator for the net's (unwanted) potential to terminate. For that, we will have a brief look at the research area of *inconsistency measurement* [9], [12], [13]. Inconsistency measurement is the field that is concerned with measuring conflicts in logic-based knowledge representation formalisms. We can identify analogies between Petri Nets and logic-based formalisms, and then apply measures similar to the existing ones, instead of developing completely new measures specifically for Petri Nets. Additionally, we will benefit from the work in the area of inconsistency measurement when we come to *rationality postulates* to assess the quality of the measures.

Approaches to inconsistency measurement exist for many different logics, but we will focus on classical propositional logic here. Propositional logic is based on *propositional variables* that can be either *true* or *false*. They can form *formulas* through connectors $\neg$ (negation), $\wedge$ (conjunction), and $\vee$ (disjunction). A set of formulas can define a *propositional language*. The assignment of a value to all propositions of a formula or set of formulas is called *interpretation*.

We will denote a knowledge base, i. e., a set of formulas, by K. With $At(X)$, where X is a formula or set of formulas, we denote the set of propositions contained in X.

A formula or set of formulas X is inconsistent iff there is no interpretation I on $At(X)$ such that every formula is satisfied by I under the classical semantics of proposition logic. Otherwise, X is consistent.

Example 1. *Consider the set of formulas K_1.*

$$K_1 = \{r \to c, s \to c, s \to f, f \to \neg r, r \wedge s\}$$

There is no interpretation, where all formulas of K_1 are satisfied, because $r \wedge s$ would need to be true. But from s follows f and from f follows $\neg r$ that cannot be true at the same time as r.

A set $M \subseteq K$ is a *minimal inconsistent subset* if M is inconsistent and every proper subset of M is consistent. For example, the set $\{s \to f, f \to \neg r, r \wedge s\}$ constitutes a minimal inconsistent subset of K_1 from Example 1. The set of all minimal inconsistent subsets of a knowledge base K is denoted by $\mathbb{M}(K)$.

A formula ϕ is called *free formula* if it is not contained in any minimal inconsistent subset, for instance $r \to c$ in Example 1.

An inconsistency measure is a function that assigns a non-negative real number to a given knowledge base. Usually, the value 0 means there is no inconsistency at all, while larger numbers indicate more inconsistency. We will have a look at a few measures based on minimal inconsistent subsets from [14], [15]. A broader overview on inconsistency measurement can be found in [9].

Definition 6. *The* $\mathbb{M}$-inconsistency measure *[14]:*

$$I_{\mathbb{M}}(K) = |\mathbb{M}(K)|.$$

The $\mathbb{M}$-inconsistency measure measures the number of minimal inconsistent subsets of the knowledge base.

Definition 7. *The* $\mathbb{M}^C$-inconsistency measure *[14]:*

$$I_{\mathbb{M}^C}(K) = \sum_{M \in \mathbb{M}(K)} \frac{1}{|M|}.$$

The $\mathbb{M}^C$-inconsistency measure measures the size of the minimal inconsistent subsets (where larger sets mean less inconsistency).

Definition 8. *The* mv-inconsistency measure *[15]:*

$$I_{mv}(K) = \frac{|\bigcup_{M \in \mathbb{M}(K)} At(M)|}{|At(K)|}.$$

The mv-inconsistency measure measures the proportion of propositions appearing in minimal inconsistent subsets.

Example 2. *We will now apply the measures to $K_1 = \{r \to c, s \to c, s \to f, f \to \neg r, r \wedge s\}$ from Example 1:*

For the $\mathbb{M}$-measure we need the set of all minimal inconsistent subsets of $K_1 : \mathbb{M}(K_1) = \{\{s \to f, f \to \neg r, r \wedge s\}\}$. The value of the $\mathbb{M}$-measure is:

$$I_{\mathbb{M}}(K_1) = |\{\{s \to f, f \to \neg r, r \wedge s\}\}| = 1.$$

The value of the $\mathbb{M}^C$-measure is:

$$I_{\mathbb{M}^C}(K_1) = \frac{1}{|\{s \to f, f \to \neg r, r \wedge s\}|} = \frac{1}{3}.$$

For the mv-measure we need the propositions contained in the minimal inconsistent set(s) of $K_1 : \bigcup_{M \in \mathbb{M}(K_1)} At(M) = At(\{s \to f, f \to \neg r, r \wedge s\}) = \{r, s, f\}$ and the set of all propositions of $K_1 : At(K_1) = \{r, c, s, f\}$. The value of the mv-measure is:

$$I_{mv}(K_1) = \frac{|\{r, s, f\}|}{|\{r, c, s, f\}|} = \frac{3}{4}.$$

Inconsistency measures are usually evaluated wrt. *rationality postulates* [16], i. e., properties describing a desirable behaviour for special scenarios. We recall some simple ones in the following.

Consistency
$I(K) = 0$ iff K consistent, i.e. the value 0 indicates the absence of inconsistency.

Normalization
$0 \leq I(K) \leq 1$. Normalization allows the comparison of the values from knowledge bases of different sizes.

Monotony
$K \subseteq K' \implies I(K) \leq I(K')$ means that adding information to the knowledge base cannot reduce inconsistency.

Free-formula independence
$I(K) = I(K \setminus \{\alpha\})$ if α is a free formula. This means removing free formulas does not change the amount of inconsistency.

Consistency is satisfied by all introduced inconsistency measures. Normalization is only satisfied by the mv-measure. Monotony and free-formula independence are satisfied by all introduced inconsistency measures, except for the mv-measure. See [16] for more discussion.

III. Deadlock-based Measures and Rationality Postulates for Petri Nets

In this section, we will define measures that aim to evaluate a net's potential to terminate, using the presence of deadlocks as an indicator. For this purpose, we will adapt existing inconsistency measures to Petri Nets and define rationality postulates. Our analogy to inconsistency of a knowledge base is a Petri Net's potential to encounter a dead marking. We use minimal deadlocks as an analogy to minimal inconsistent subsets, since both are the smallest subsets which correspond with the conflict resp. unwanted behaviour.

A. Measures

We begin by applying the three measures for classical propositional logic introduced in Section II-C to minimal deadlocks instead of minimal inconsistent subsets. We differentiate between two types of measures: Measures that take tokens into account and those that do not.

1) Token-dependent Measures: Token-dependent measures aim to assess the net's potential to terminate for a given (initial) marking M. Since deadlocks containing a marked trap do not contribute to the net's potential to terminate, we will use measures based on minimal deadlocks without marked traps.

Recall the definition of a Petri Net $N = (P, T, I, O)$, where P is the set of all places of the net, T is the set of all transitions of the net, I is the set of input edges, and O is the set of output edges. Recall also the notation (N, M) for a marked net and $\mathbb{D}(N)$ for the set of all minimal deadlocks of a net N.

Additionally we define $\mathbb{D}_M(N)$, the set of minimal deadlocks $D_M \in \mathbb{D}(N)$ of the net (N, M), such that $D'_M \subseteq D_M \implies D'^{\bullet}_M \not\subseteq {}^{\bullet}D'_M$ or $\forall\ p \in D'_M : M(p) = 0$. In other words, $\mathbb{D}_M(N)$ is the set of all minimal deadlocks of N without a trap that is marked in M.

Token-dependent measures are then functions taking a net N and a marking M and returning a non-negative real value $T(N, M)$.

Definition 9. *We define the $\mathbb{D}_M$-measure as:*

$$T_{\mathbb{D}_M}(N, M) = |\mathbb{D}_M(N)|.$$

The $\mathbb{D}_M$-measure measures the number of minimal deadlocks that do not contain a trap that is marked at the marking M. The value 0 means there are no (minimal) deadlocks without a marked trap.

Definition 10. *We define the $\mathbb{D}^C_M$-measure as:*

$$T_{\mathbb{D}^C_M}(N, M) = \sum_{D_M \in \mathbb{D}_M(N)} \frac{1}{|D_M|}.$$

The $\mathbb{D}^C_M$-measure measures the size of the minimal deadlocks without initially (at marking M) marked trap. The larger the minimal deadlocks without marked trap, the smaller the value of this measure. This means, larger deadlocks without marked trap are considered less problematic. Note that, due to the empty sum evaluating to zero, we have $T_{\mathbb{D}^C_M}(N, M) = 0$ if $\mathbb{D}_M(N) = \emptyset$.

Definition 11. *We define the mr_M-measure as:*

$$T_{mr_M}(N, M) = \frac{|\bigcup_{D_M \in \mathbb{D}_M(N)} D_M|}{|P|}.$$

The mr_M-measure measures the proportion of places appearing in minimal deadlocks without (at marking M) marked trap.

Example 3. *We will now calculate the values of the introduced measures for the marked nets (N, M_1) shown in Figure 14 and (N, M_2) shown in Figure 15. To spot deadlocks and traps manually, we use the intuition, that a deadlock either contains a loop, or a place with empty pre-set, and a trap either contains a loop, or a place with empty post-set. We start by searching for deadlocks, since traps are only relevant when they are inside a minimal deadlock.*

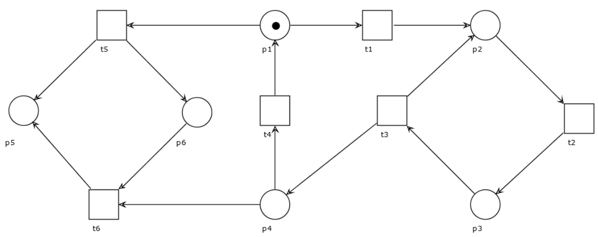

Figure 14. A marked Petri Net

In Figure 14, we find $D_1 = \{p1, p2, p3, p4\}$ *with* ${}^\bullet D_1 = \{t1, t2, t3, t4\} \subseteq D_1{}^\bullet = \{t1, t2, t3, t4, t5, t6\}$ *as only (and thus minimal) deadlock. Within this deadlock we can find trap* $S_1 = \{p2, p3\}$*, since* $S_1{}^\bullet = \{t2, t3\} \subseteq {}^\bullet S_1 = \{t1, t2, t3\}$.

In Figure 14, S_1 *is unmarked, thus for the marked net* (N, M_1)*:*

$T_{\mathbb{D}_M}(N, M_1) = |\mathbb{D}_{M_1}(N)| = |\{D_1\}| = 1$

$T_{\mathbb{D}_M^C}(N, M_1) = \sum_{D_M \in \mathbb{D}_{M_1}(N)} \frac{1}{|D_M|} = \frac{1}{|D_1|} = \frac{1}{4}$

$T_{mr_M}(N, M_1) = \frac{|\bigcup_{D_M \in \mathbb{D}_{M_1}} D_M|}{|P|} = \frac{|D_1|}{|P|} = \frac{2}{3}$

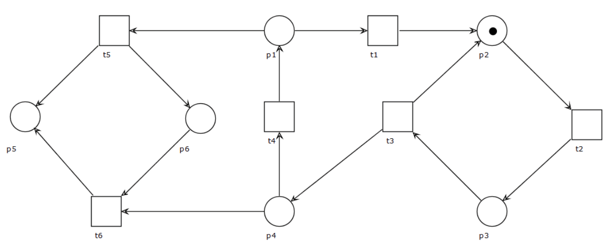

Figure 15. The same net with different marking

In Figure 15 the net is the same as in Figure 14, except for the marking. Thus we have minimal deadlock D_1 *and the contained trap* S_1 *again.*

In Figure 15, S_1 *is marked, therefore for the marked net* (N, M_2)*:*

$T_{\mathbb{D}_M}(N, M_2) = |\mathbb{D}_{M_2}(N)| = |\emptyset| = 0$

$T_{\mathbb{D}_M^C}(N, M_2) = \sum_{D_M \in \mathbb{D}_{M_2}(N)} \frac{1}{|D_M|} = 0$

$T_{mr_M}(N, M_2) = \frac{|\bigcup_{D_M \in \mathbb{D}_{M_2}} D_M|}{|P|} = \frac{|\emptyset|}{|P|} = 0$

2) Token-independent Measures: Token-independent measures aim to assess a net's potential to terminate regardless of the initial marking. The three measures we will introduce here are based on the minimal deadlocks of the net, disregarding whether they contain a marked trap or not. Therefore the token-independent measures can be interpreted as a measure for the worst markings where all deadlocks are unmarked.

Token-independent measures are then functions taking a net N and returning a non-negative real value $T(N)$.

Definition 12. *We define the* $\mathbb{D}$-measure *as:*

$$T_{\mathbb{D}}(N) = T_{\mathbb{D}_M}(N, M_\varepsilon) = |\mathbb{D}_{M_\varepsilon}(N)| = |\mathbb{D}(N)|,$$

where M_ε *is the marking at which all places of the net are unmarked.*

At marking M_ε, the set of minimal deadlocks without marked trap $\mathbb{D}_{M_\varepsilon}(N)$ equals the set of all minimal deadlocks $\mathbb{D}(N)$, since all deadlocks are unmarked. The $\mathbb{D}$-measure measures the number of minimal deadlocks of the net. 0 means there are no minimal deadlocks (and therefore no deadlocks at all).

Definition 13. *We define the* $\mathbb{D}^C$-measure *as:*

$$T_{\mathbb{D}^C}(N) = T_{\mathbb{D}_M^C}(N, M_\varepsilon) = \sum_{D_M \in \mathbb{D}_{M_\varepsilon}(N)} \frac{1}{|D_M|} = \sum_{D \in \mathbb{D}(N)} \frac{1}{|D|},$$

where M_ε *is the marking at which all places of the net are unmarked.*

At marking M_ε, the set of minimal deadlocks without marked trap $\mathbb{D}_{M_\varepsilon}(N)$ equals the set of all minimal deadlocks $\mathbb{D}(N)$, since all deadlocks are unmarked.

The $\mathbb{D}^C$-measure measures the size of the minimal deadlocks. The larger the minimal deadlocks, the smaller the value of this measure. This means, larger deadlocks are considered less problematic. The idea is, that larger deadlocks—as larger inconsistent subsets—are more "hidden" or indirect than smaller ones. Therefore they could be less likely to get empty or to be unmarked initially.

Definition 14. *We define the* mr-measure *as:*

$$T_{mr}(N) = T_{mr_M}(N, M_\varepsilon) = \frac{|\bigcup_{D_M \in \mathbb{D}_{M_\varepsilon}(N)} D_M|}{|P|} = \frac{|\bigcup_{D \in \mathbb{D}(N)} D|}{|P|},$$

where M_ε *is the marking at which all places of the net are unmarked.*

At marking M_ε, the set of minimal deadlocks without marked trap $\mathbb{D}_{M_\varepsilon}(N)$ equals the set of all minimal deadlocks $\mathbb{D}(N)$, since all deadlocks are unmarked.

The mr-measure measures the proportion of places appearing in minimal deadlocks.

Example 4. *We will now calculate the values of all six measures – token-dependent and token-independent ones – for the net below (Figure 16). We will identify and count the minimal deadlocks and the minimal deadlocks without marked trap.*

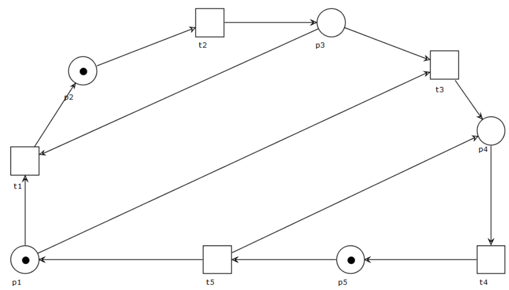

Figure 16. A marked Petri Net

The minimal deadlocks of the depicted net (N, M_1) *are* $D_1 = \{p2, p3\}$ *and* $D_2 = \{p1, p4, p5\}$. *The deadlock* D_1 *does not contain a trap. The deadlock* D_2 *contains the trap* $S_1 = \{p4, p5\}$, *which is marked at the depicted marking.*

$$T_{\mathbb{D}}(N) = |\{D_1, D_2\}| = 2$$

$$T_{\mathbb{D}_M}(N, M_1) = |\{D_1\}| = 1$$

$$T_{\mathbb{D}^C}(N) = \frac{1}{|D_1|} + \frac{1}{|D_2|} = \frac{1}{2} + \frac{1}{3} = \frac{5}{6}$$

$$T_{\mathbb{D}_M^C}(N, M_1) = \frac{1}{|D_1|} = \frac{1}{2}$$

$$T_{mr}(N) = \frac{|D_1 \cup D_2|}{|P|} = \frac{|\{p1,p2,p3,p4,p5\}|}{|\{p1,p2,p3,p4,p5\}|} = \frac{5}{5} = 1$$

$$T_{mr_M}(N, M_1) = \frac{|D_1|}{|P|} = \frac{|\{p2,p3\}|}{|\{p1,p2,p3,p4,p5\}|} = \frac{2}{5}$$

As we can observe in Example 3 and 4, the values of the different measures can differ a lot from each other for the same net. Therefore, it is not clear what the value of a single measure means for the termination potential of a concrete net in detail. However, the use of a combination of multiple measures should yield a rough idea about the condition of the analyzed net. The token-dependent measures should be viewed as more accurate, since they are only taking the minimal deadlocks relevant to the termination potential into account.

B. Rationality Postulates

In this section we will define some rationality postulates for Petri Nets and show which measures satisfy them. Formally we will define the postulates for token-dependent measures, because the corresponding token-independent measures are just a special case of them. Therefore, for the remainder of this section let T be an arbitrary token-dependent measure.

Our first postulate in the context of Petri Nets is *deadlock-freeness* and it characterises the meaning of the minimal value of a measure. We say T satisfies *deadlock freeness* iff

$$T(N, M) = 0 \implies (N, M) \text{ is deadlock-free.}$$

For token-independent measures the value 0 means the net is deadlock-free for all initial markings, i.e. there is no dead marking reachable from any initial marking.[2] If the value of a measure is different from 0, the net is not necessarily not deadlock-free, i.e. the net could still be deadlock-free. We define deadlock freeness in this manner, because otherwise all measures based on deadlocks would automatically fail to satisfy it, since a net can be deadlock-free, although it contains deadlocks. It could, for example, contain a source transition, thus making a dead marking impossible (see Figure 17). Even without a source transition, a net could be deadlock-free for some initial marking(s), if the structure of the contained deadlocks and their neighbourhood prevent the deadlocks from getting empty, although some of the deadlocks contain no trap (see Example 5).

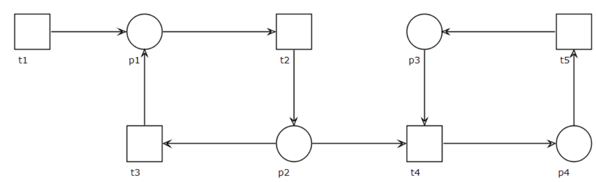

Figure 17. Deadlock-free Petri Net despite contained deadlock (without marked trap)

Example 5. *In Figure 18, the transition* $t5$ *is a sink transition, and* $D_1 = \{p1, p2, p3, t4\}$ *is a deadlock, since* ${}^\bullet D_1 = \{t1, t2, t3, t4\} \subseteq D_1{}^\bullet = \{t1, t2, t3, t4, t5\}$. D_1 *can only get empty if* $t5$ *is fired, but since* $p5$ *is not marked,* $t5$ *cannot fire until* $t1$ *got enabled and then fired. When* $t1$ *is fired, there are two tokens produced inside the deadlock, but only one in* $p5$, *thus* $t5$ *can only fire once. Therefore, at least one token stays inside the deadlock. Since this can happen infinitely often, the net is live and deadlock-free.*

If $p4$ *and* $p5$ *would both contain exactly one token at the initial marking (all other places empty),* $t5$ *could consume both tokens, thus resulting in a dead marking.*

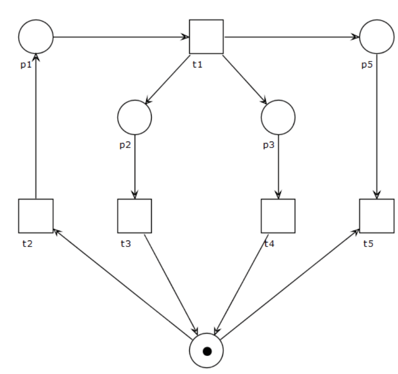

Figure 18. Deadlock-free (and live) marked Petri Net despite contained deadlock without any trap; suggested by Li Jiao et al. [17]

In order to assess which measures fulfill deadlock freeness, we need to understand the connection between dead markings and deadlocks. As mentioned in Section II-B, if there are no deadlocks without initially marked trap, there is no dead marking reachable [4]. This means, without the existence of an

[2]In the terms deadlock-free and deadlock freeness, deadlock stands for dead marking, but apart from that, in this work, the term deadlock indicates the structure from Definition 3.

unmarked deadlock, the marking cannot be dead. And since a marked trap can never get empty (see Section II-B), a deadlock containing it cannot get empty as well. Therefore a marked net is deadlock-free if it does not contain deadlocks without marked trap. This means the absence of deadlocks (without marked trap) implies deadlock freeness. Therefore, whenever the value 0 of a measure indicates the absence of deadlocks (without marked trap), the measure fullfills deadlock freeness.

Our next property is about monotony. In the context of inconsistency measurement in classical propositional logic, it is clear that inconsistency cannot be resolved by adding formulas to a knowledge base. Therefore, a widely accepted property of inconsistency measures [16] is that they should behave monotonic, i. e., adding formulas cannot decrease the inconsistency value. For Petri Nets the situation is a bit different, since extending a net may resolve deadlocks. We will investigate this issue a bit deeper now.

Adding or removing places or transitions to a net without connecting them by edges will not change the existing deadlocks and traps, since they are determined by their pre-set and post-set, which depends on edges, and thus cannot change by adding or removing unconnected places or transitions. Unconnected places constitute minimal deadlocks themselves, and they can be added to a deadlock and the resulting set of places will be a deadlock as well. Therefore, the number of deadlocks and minimal deadlocks will increase. However, these unconnected places do not have any influence on the net's behaviour.

Adding and removing edges results in different pre-sets or post-sets of places. Therefore, this can affect the property of a set of places being a deadlock or trap.

Adding input edges means that the connected transitions are added to the pre-set of the places they are connected to. Let P be a set of places and $e = (p, t) \in I$ an edge, such that $p \in P$, $t \notin {}^\bullet P$ and ${}^\bullet P \subseteq P^\bullet$. Adding edge e means ${}^\bullet P_{new} = {}^\bullet P \cup \{t\}$. If $t \notin P^\bullet$, then ${}^\bullet P_{new} \nsubseteq P^\bullet$, so P is not a deadlock any more. It was resolved by adding e. Analogous, a deadlock can emerge by removing an input edge. Thus adding input edges can resolve deadlocks (see Figure 19 and 20) and removing input edges can generate deadlocks.

Adding output edges means that the connected transitions are added to the post-set of the places they are connected to. Let P be a set of places and $e = (p, t) \in O$ an edge, such that $p \in P$, $t \notin P^\bullet$ and ${}^\bullet P \nsubseteq P^\bullet$. Adding edge e means $P^\bullet_{new} = P^\bullet \cup \{t\}$. If ${}^\bullet P \setminus P^\bullet = \{t\}$, then ${}^\bullet P \subseteq P^\bullet_{new}$. This means P is a deadlock after adding edge e, but was not before. Analogous, a deadlock can be resolved by removing an output edge. Thus adding output edges can generate deadlocks (see Figure 21 and 22) and removing output edges can resolve deadlocks.

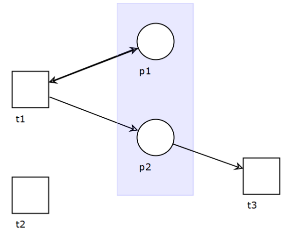

Figure 19. Deadlock

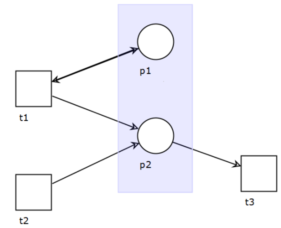

Figure 20. Deadlock resolved by adding an input edge

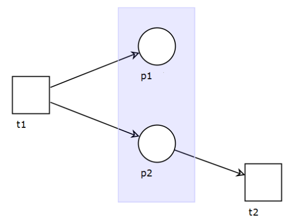

Figure 21. No deadlock

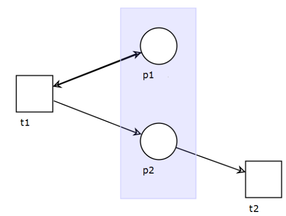

Figure 22. New deadlock created by adding an output edge

Table I shows, case by case, how sets of places can or cannot change their property of being a deadlock or trap by adding and removing input and output edges. P denotes the set of places we are investigating. ${}^\bullet P_{new}$ denotes the pre-set of P after adding or removing an input edge connected to a place of P. Analogous, $P^\bullet{}_{new}$ denotes the post-set of P after adding or removing an output edge. t denotes a transition that is added or removed from the pre-set or post-set of P because of the corresponding added or removed edge. For adding/removing input edges, t is added/removed to/from the pre-set of P and for output edges it is added/removed to/from the post-set of P.

From the properties described above, we obtain a monotonic behaviour (regarding the number of minimal deadlocks and minimal deadlocks without marked trap) for adding output edges and for removing input edges. Therefore, we say a measure T satisfies *weak monotony*, if the value is equal or higher when output edges are added or input edges are removed, formally

$$O_N \subseteq O_{N'}, I_N \supseteq I_{N'} \implies T(N, M) \leq T(N', M).$$

We now turn to the aspect of decomposability, i. e., how do the values of our measures aggregate when nets are combined? If adding two nets without connecting them to each other (and for a marked net, without changing the distribution of the tokens) results in the values of the measure of the two nets adding up as well, the measure satisfies *weak additivity*. Formally, we say T satisfies *weak additivity* iff

$$P_N \cap P_{N'} = T_N \cap T_{N'} = \emptyset \implies T((N, M_N) + (N', M_{N'})) = T(N, M_N) + T(N', M_{N'})$$

	P is a deadlock $^\bullet P \subseteq P^\bullet$	P is not a deadlock $^\bullet P \not\subseteq P^\bullet$	P is a trap $P^\bullet \subseteq {}^\bullet P$	P is not a trap $P^\bullet \not\subseteq {}^\bullet P$	
Adding an input edge $^\bullet P \subseteq {}^\bullet P_{new}$	*for* $t \notin P^\bullet$: $^\bullet P_{new} \not\subseteq P^\bullet$	$^\bullet P_{new} \not\subseteq P^\bullet$	$P^\bullet \subseteq {}^\bullet P_{new}$	*for* $t = P^\bullet \setminus {}^\bullet P$: $P^\bullet \subseteq {}^\bullet P_{new}$	potentially fewer deadlocks and more traps
Removing an input edge $^\bullet P_{new} \subseteq {}^\bullet P$	$^\bullet P_{new} \subseteq P^\bullet$	*for* $t = {}^\bullet P \setminus P^\bullet$: $^\bullet P_{new} \subseteq P^\bullet$	*for* $t \in P^\bullet$: $P^\bullet \not\subseteq {}^\bullet P_{new}$	$P^\bullet \not\subseteq {}^\bullet P_{new}$	potentially more deadlocks and fewer traps
Adding an output edge $P^\bullet \subseteq P^\bullet_{new}$	$^\bullet P \subseteq P^\bullet_{new}$	*for* $t = {}^\bullet P \setminus P^\bullet$: $^\bullet P \subseteq P^\bullet_{new}$	*for* $t \notin {}^\bullet P$: $P^\bullet_{new} \not\subseteq {}^\bullet P$	$P^\bullet_{new} \not\subseteq {}^\bullet P$	
Removing an output edge $P^\bullet_{new} \subseteq P^\bullet$	*for* $t \in {}^\bullet P$: $^\bullet P \not\subseteq P^\bullet_{new}$	$^\bullet P \not\subseteq P^\bullet_{new}$	$P^\bullet_{new} \subseteq {}^\bullet P$	*for* $t = P^\bullet \setminus {}^\bullet P$: $P^\bullet_{new} \subseteq {}^\bullet P$	potentially fewer deadlocks and more traps

Table I
CONSEQUENCES OF ADDING AND REMOVING INPUT AND OUTPUT EDGES

where $(N, M_N) + (N', M_{N'}) = (N + N', M_{N+N'})$ *with* $N + N' = (P_N \cup P_{N'}, T_N \cup T_{N'}, I_N \cup I_{N'}, O_N \cup O_{N'})$ *and*

$$M_{N+N'} = \begin{pmatrix} M_N(p_{N1}) \\ \vdots \\ M_N(p_{Nn}) \\ M_{N'}(p_{N'1}) \\ \vdots \\ M_{N'}(p_{N'm}) \end{pmatrix} \qquad \textit{for } |P_N| = n \textit{ and}$$

$|P_{N'}| = m.$

If we add two nets without connecting them to each other, the number of minimal deadlocks adds up as well, i. e. it will not change for either part of the net. This is, because without removing or adding new edges, existing deadlocks cannot change. But since the union of deadlocks is a deadlock as well, the number of non-minimal deadlocks can change.

Our final property is concerned with the range of possible values of a measure. We say T satisfies *Normalization* iff

$$0 \leq T(N, M) \leq 1.$$

The above four postulates constitute basic desirable properties for our measures. Part of ongoing work is to adapt further postulates from, e. g., [16].

Table II shows the compliance of our measures with the above postulates. The formal proofs of these statements can be found in an online appendix[3].

[3] http://mthimm.de/misc/proofs_cbi21_udt.pdf

	Deadlock Freeness	Weak Monotony	Weak Additivity	Normalization
$\mathbb{D}_M$, $\mathbb{D}$	✔	✔	✔	✘
$\mathbb{D}^C_M$, $\mathbb{D}^C$	✔	✔	✔	✘
mr_M, mr	✔	✘	✘	✔

Table II
MEASURES AND POSTULATES

IV. DISCUSSION

The topic of this paper is combining the fields of inconsistency measurement and liveness and deadlock analysis in Petri Nets. The field of inconsistency measurement served as an inspiration and a template for the measures and postulates introduced in this paper, while in substance, the topic of this paper is closer to liveness analysis.

To the best of our knowledge, this is the first approach to quantitatively assess a Petri Net's potential to terminate, beyond stating whether it is deadlock-free or not. However, there has been a lot of research done around deadlocks and minimal deadlocks, in particular how to find them and how to prevent them from getting empty, see e. g. [6], [7].

In this work, novel evaluation criteria for ordinary marked Petri Net models have been developed. For this purpose, the parallels between minimal inconsistent sets in propositional logic and minimal deadlocks have been exploited, as they are both the smallest subsets that can be associated with the conflict/unwanted behaviour. In this work, three inconsistency measures for minimal inconsistent sets have been adapted to Petri Nets. By defining the three measures in two different ways, six distinct measures for Petri Nets have been obtained. Furthermore, four rationality postulates for Petri Net measures have been defined. The measures can be used to reason about the reliability of processes or systems modeled with ordinary marked Petri Nets, without the need to compute reachable markings.

V. Future Work

In this work, we focused on ordinary marked nets, but there are several other types of Petri Nets, for example free-choice nets, bounded nets, nets with multi-edges (edge weights) and more (see for example [2], [4], [5]). Free-choice nets are only forward branching, thus often making the analysis easier. Bounded nets have limits for the amount of tokens per place, so-called place capacities, creating another cause for not firable transitions because of insufficient place capacities in the post-set. It would be interesting to develop measures for other types of nets as well and to investigate whether the measures introduced in this work can be transferred to different types of Petri Nets.

In the field of deadlock prevention (here: prevention of dead markings), often the initial marking gets chosen such that there are no deadlocks without marked trap – if possible (see [7] for example). The information about how many places would need to be marked for this purpose could serve as another measure for termination potential. Interesting are also approaches based on elementary deadlocks [6], where the strict minimal deadlocks (i. e. minimal deadlocks that do not contain a trap) are categorized as elementary and dependent ones, such that it is sufficient to monitor the elementary deadlocks. This means, it might be possible to reduce the problem of deadlock freeness / termination to less than the set of all minimal deadlocks without marked trap. It seems therefore promising to develop further measures specifically for Petri Nets inspired by concepts used in deadlock prevention.

Another interesting aspect is the relation between deadlock freeness and liveness, that has already been investigated and is equivalent for some types of nets and under special restrictions (see for example [18] and [19]). Although a non-live but deadlock-free net is considered less problematic than a net that is not deadlock-free (and therefore not live by definition), liveness is also a desirable property for Petri Nets. As for deadlock freeness, there has been some research already. For example, it has been found, that liveness of an ordinary marked Petri Net is preserved when adding additional tokens to the initial marking, iff every deadlock contains at least one trap (this property is called *Liveness Monotonicity* [17]). It seems useful to develop measures for liveness as well.

Acknowledgements

This research is part of the research project "Handling Inconsistencies in Business Process Modeling", which is funded by the German Research Association (project number 424710479).

References

[1] T. Agerwala, "Special feature: Putting petri nets to work," *Computer*, vol. 12, no. 12, pp. 85–94, 1979.

[2] B. Baumgarten, *Petri-Netze. Grundlagen und Anwendung*. Mannheim: Bibliographisches Institut & F.A. Brockhaus AG, 1990.

[3] L. P. und H. Wimmel, *Theoretische Informatik. Petri-Netze*. Berlin: Springer-Verlag, 2003.

[4] W. Reisig, *Petrinetze. Eine Einführung*. Berlin: Springer-Verlag, 1986.

[5] P. H. Starke, *Petri-Netze*. Berlin: VEB Deutscher Verlag der Wissenschaften, 1980.

[6] Z. Li and M. Zhou, *Deadlock Control Based on Elementary Siphons*. London: Springer, 2009, pp. 107–157.

[7] Y. Hou and K. Barkaoui, "Deadlock analysis and control based on petri nets: A siphon approach review," *Advances in Mechanical Engineering*, vol. 9, no. 5, pp. 1–30, 2017.

[8] J. L. Peterson, *Petri Net Theory and the Modeling of Systems*, 1981-2019.

[9] M. Thimm, "Inconsistency measurement," in *Proceedings of the 13th International Conference on Scalable Uncertainty Management (SUM'19)*, ser. Lecture Notes in Artificial Intelligence, N. B. Amor, B. Quost, and M. Theobald, Eds., vol. 11940. Springer International Publishing, December 2019, pp. 9–23.

[10] J. Wang, *Timed Petri Nets. Theory and Application*. Norwell: Kluwer Academic Publishers, 1998.

[11] M. Minoux and K. Barakaoui, "Deadlocks and traps in petri nets as horn-satisfiability solutions and some related polynomially solvable problems," *Discrete Applied Mathematics*, vol. 29, pp. 195–210, 1990.

[12] J. Grant, "Classifications for inconsistent theories," *Notre Dame Journal of Formal Logic*, vol. 19, no. 3, pp. 435–444, 1978.

[13] J. Grant and M. V. Martinez, Eds., *Measuring Inconsistency in Information*, ser. Studies in Logic. College Publications, 2018, vol. 73.

[14] A. Hunter and S. Konieczny, "Measuring inconsistency through minimal inconsistent sets," in *Proceedings of the Eleventh International Conference on Principles of Knowledge Representation and Reasoning (KR'2008)*, 2008, pp. 358–366.

[15] G. Xiao and Y. Ma, "Inconsistency measurement based on variables in minimal unsatisfiable subsets," in *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI'12)*, 2012.

[16] M. Thimm, "On the evaluation of inconsistency measures," in *Measuring Inconsistency in Information*, ser. Studies in Logic, J. Grant and M. V. Martinez, Eds. College Publications, February 2018, vol. 73.

[17] L. Jiao, T. Cheung, and W. Lu, "Characterizing liveness of petri nets in terms of siphons," in *Application and Theory of Petri Nets 2002*, J. Esparza and C. Lakos, Eds. Berlin: Springer-Verlag, 2002, pp. 203–216.

[18] K. Barkaoui, J. Couvreur, and K. Klai, "On the equivalence between liveness and deadlock-freeness in petri nets," in *Applications and Theory of Petri Nets 2005*. Berlin: Springer-Verlag, 2005, pp. 90–107.

[19] E. Teruel and M. Silva, "Structure theory of equal conflict systems," *Theoretical Computer Science*, vol. 153, no. 1, pp. 271 – 300, 1996.