

# Knowledge Representation in Probabilistic Spatio-Temporal Knowledge Bases

**Francesco Parisi**

*Department of Informatics, Modeling,  
Electronics and System Engineering  
University of Calabria, Rende, Italy*

FPARISI@DIMES.UNICAL.IT

**John Grant**

*Department of Computer Science and UMIACS  
University of Maryland, College Park, USA*

GRANT@CS.UMD.EDU

## Abstract

We represent knowledge as integrity constraints in a formalization of probabilistic spatio-temporal knowledge bases. We start by defining the syntax and semantics of a formalization called PST knowledge bases. This definition generalizes an earlier version, called SPOT, which is a declarative framework for the representation and processing of probabilistic spatio-temporal data where probability is represented as an interval because the exact value is unknown. We augment the previous definition by adding a type of non-atomic formula that expresses integrity constraints. The result is a highly expressive formalism for knowledge representation dealing with probabilistic spatio-temporal data. We obtain complexity results both for checking the consistency of PST knowledge bases and for answering queries in PST knowledge bases, and also specify tractable cases. All the domains in the PST framework are finite, but we extend our results also to arbitrarily large finite domains.

## 1. Introduction

Recent years have seen a great deal of interest in tracking moving objects. This is a fundamental issue for many applications providing location-based and context-aware services, such as emergency call-out assistance, live traffic reports, food and drink finder, location-based advertising, mobile tourist guidance, pervasive healthcare, and analysis of animal behavior (Ahson & Ilyas, 2010; Petrova & Wang, 2011; Karimi, 2013). Such innovative services are becoming so widely diffused that MarketsandMarkets forecasts that the location-based services market will grow from \$8.12 billion in 2014 to \$39.87 billion in 2019 (MarketsandMarkets, 2014).

An important aspect of the systems providing location-based and context-aware services is that they need to manage spatial and temporal data together. For this reason, researchers have investigated in detail the representation and processing of spatio-temporal data, both in AI (Cohn & Hazarika, 2001; Gabelaia, Kontchakov, Kurucz, Wolter, & Zakharyashev, 2005; Yaman, Nau, & Subrahmanian, 2004, 2005a; Knapp, Merz, Wirsing, & Zappe, 2006) and databases (Agarwal, Arge, & Erickson, 2003; Pelanis, Saltenis, & Jensen, 2006). However, in many cases the location of objects is uncertain: such cases can be handled by using probabilities (Parker, Yaman, Nau, & Subrahmanian, 2007b; Tao, Cheng, Xiao, Ngai, Kao, & Prabhakar, 2005). Sometimes the probabilities themselves are not known exactly. Indeed, the position of an object at a given time is estimated by means of a location estimation method such as proximity (where the location of an object is derived from its vicinity to one or more antennas), fingerprinting (where radio signal strength mea-

surements produced by a moving object are matched against a radio map that is built before the system is working), and dead reckoning (where the position of an object is derived from the last known position, assuming that the direction of motion and either the speed or the travelled distance are known) (Ahson & Ilyas, 2010; Karimi, 2013). However, since location estimation methods have limited accuracy and precision, what can be asserted is that an object is at a given position at a given time with a probability whose value belongs to an interval. The SPOT (Spatial PrObabilistic Temporal) framework was introduced by Parker, Subrahmanian, and Grant (2007a) to provide a declarative framework for the representation and processing of probabilistic spatio-temporal data with probabilities that are not known exactly.

The SPOT framework is able to represent atomic statements of the form “object *id* is/was/will be inside region *r* at time *t* with probability in the interval  $[\ell, u]$ ”. This allows the representation of information concerning moving objects in several application domains. A cell phone provider is interested in knowing which cell phones will be in the range of some towers at a given time and with what probability (Bayir, Demirbas, & Eagle, 2010). A transportation company is interested in predicting the vehicles that will be on a given road at a given time (and with what probability) in order avoid congestion (Karbassi & Barth, 2003). Finally, a retailer is interested in knowing the positions of the shoppers moving in a shopping mall in order to offer suitable customized coupons on discounts (Kurkovsky & Harihar, 2006).

The framework introduced by Parker et al. (2007a) was then extended by Parker, Infantes, Subrahmanian, and Grant (2008) and Grant, Parisi, Parker, and Subrahmanian (2010) to include the specific integrity constraint that, for a given moving object, only some points are reachable from a given starting point in one time unit. This captures the scenario where objects have speed limits and only some points are reachable by objects depending on the distance between the points. However, even such an extended SPOT framework is not yet general enough to represent additional knowledge concerning the movements of objects. Examples of knowledge we may be aware of but cannot represent in the SPOT framework are, for instance, the fact that

- (i) there cannot be two distinct objects in a given region in a given time interval (as it happens during airport passenger screening);
- (ii) some object cannot reach a given region starting from a given location in less than a given amount of time (as happens for the vehicles whose route options as well as their speed are limited);
- (iii) an object can go away from a given region only if it stayed there for at least a given amount of time (as happens in production lines where assembling several parts requires a given amount of time).

To overcome such limitation and allow this kind of knowledge to be represented, we define probabilistic spatio-temporal (PST) knowledge bases (KBs) consisting of atomic statements, such as those representable in the SPOT framework and *spatio-temporal denial (abbreviated to std) formulas*, a general class of formulas that account for all the three cases above, and many more (including the reachability constraint of Parker et al., 2008; Grant et al., 2010).

The focus of this paper is the systematic study of knowledge representation in probabilistic spatio-temporal data. We start by defining the concept of a PST KB and provide its formal semantics, which is given in terms of worlds, interpretations, and models (Section 2). We define the concept of a consistent PST KB, and characterize the complexity of checking consistency, showing

that it is NP-complete in general, and even in the presence of binary std-formulas only (Section 3). Then we present a sufficient condition for checking consistency that relies on the feasibility of a set of mixed-binary linear inequalities (Section 3.2), and a necessary condition using instead a set of linear inequalities (Section 3.3). After showing that the special case of unary std-formulas is tractable (Section 3.4), we deal with the restricted but expressive class of binary std-formulas and identify cases where the consistency checking problem is tractable (Section 3.5). We then address the problem of answering selection queries in PST KBs under both the optimistic and the cautious semantics (Section 4). We show that checking consistency can be exploited to answer such kinds of queries in PST KBs, and characterize the complexity of the query answering problem (Section 4.2). After that, we derive several sets of linear inequalities for answering queries (Section 4.3). Finally, we extend the framework to the case where time, space, and the number of objects are increased to arbitrarily large finite domains, and show that each PST KB is either eventually consistent or eventually inconsistent (Section 5). We then discuss related work (Section 6). Section 7 summarizes the paper. We also suggest research projects in Section 8.

## 2. The PST Framework

This section introduces the syntax and semantics of PST KBs generalizing the SPOT framework introduced by Parker et al. (2007a) and further extended by Parker et al. (2008) and Grant et al. (2010). Basically, we define a PST KB by augmenting the previous framework with non-atomic formulas (i.e., spatio-temporal denial formulas) that represent integrity constraints. This way we can make statements whose meaning is that certain object trajectories cannot occur.

### 2.1 Syntax

We assume the existence of three types of constant symbols: object symbols, time value symbols, and spatial region symbols. The constants are in  $ID = \{id_1, \dots, id_m\}$ ,  $T = [0, 1, \dots, tmax]$  (where  $tmax$  is an integer), and the set of  $r \subseteq Space = \{p_1, \dots, p_n\}$ . Each  $r$  is a region of  $Space$ . We apply the unique name assumption; so, for instance  $id_i$  and  $id_j$  for  $i \neq j$  are different objects; similarly,  $p_i$  and  $p_j$  for  $i \neq j$  are different points. We also use variables for each type: object variables, time variables, and spatial variables.

A *spatio-temporal atom* (*st-atom*, for short) is an expression of the form  $loc(X, Y, Z)$ , where:

- (i)  $X$  is an object variable or a constant  $id \in ID$ ,
- (ii)  $Y$  is a space<sup>1</sup> variable or a constant  $r \subseteq Space$ ,
- (iii)  $Z$  is a time variable or a constant  $t$  in  $T$ .

We say that st-atom  $loc(X, Y, Z)$  is *ground* if all of its arguments  $X, Y, Z$  are constants. For instance,  $loc(id, r, t)$ , where  $id \in ID$ ,  $r \subseteq Space$ , and  $t \in T$  is a ground st-atom. The intuitive meaning of  $loc(id, r, t)$  is that object  $id$  is/was/will be inside region  $r$  at time  $t$ .

**Definition 1** (PST atom). A PST atom is a ground st-atom  $loc(id, r, t)$  annotated with a probability interval  $[\ell, u] \subseteq [0, 1]$  (with both  $\ell$  and  $u$  rational numbers), and denoted as  $loc(id, r, t)[\ell, u]$ .

1. We write  $Space$  to refer to the set of points used in a PST KB. We write “space” to refer to the spatial aspect of probabilistic spatio-temporal knowledge.

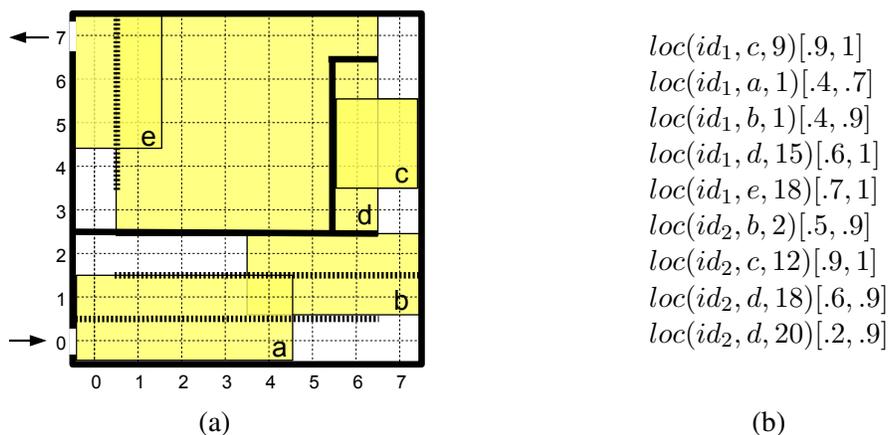


Figure 1: (a) A map of an airport area (names of regions are on their bottom-right corner); (b) PST atoms.

Intuitively, the PST atom  $loc(id, r, t)[\ell, u]$  says that object  $id$  is/was/will be inside region  $r$  at time  $t$  with probability in the interval  $[\ell, u]$ . Hence, PST atoms can represent information about the past and the present, but also information about the future, such as from methods for predicting the destination of moving objects (Mittu & Ross, 2003; Hammel, Rogers, & Yetso, 2003; Southey, Loh, & Wilkinson, 2007), or from querying predictive databases (Akdere, Cetintemel, Riondato, Upfal, & Zdonik, 2011; Parisi, Sliva, & Subrahmanian, 2013).

In the original SPOT definition, for ease of implementation, *Space* was a grid within which only rectangular regions were considered; however, in our general framework, *Space* is arbitrary and a region is any nonempty subset of *Space*. Still, for convenience we use such rectangular regions in our running example.

**Example 1.** Consider an airport security system which collects data from biometric sensors as well as from Bluetooth or WiFi enabled devices. Biometric data such as faces recognized by sensors (Li & Jain, 2011) are matched against given profiles (such as those of checked-in passports, or of wanted criminals). Similarly, device identifiers (e.g., MAC addresses) recognized in the areas covered by network antennas are matched against profiles collected by the airport hotspots (such as logins, possibly associated with passport numbers). A simplified plan of an airport area is reported in Figure 1(a), where regions  $a, b, c, d, e$  covered by sensors and/or antennas are highlighted. Once entered in this area, passengers typically move through the path delimited by queue dividers (represented by dotted lines in the figure, and overlapping with regions  $a$  and  $b$ ), and reach the room on the upper-half right side where security checks are performed (region  $c$  is included in this room). Next, passengers can spend some time in the hall room (overlapping with region  $d$ ), and finally go towards the exit (near region  $e$ ).

Suppose that the security system uses the SPOT framework to represent the information where every PST atom consists of the profile  $id$  resulting from the matching phase, the region where the sensor/antenna recognizing the profile is operating, the time point at which the profile is recognized, and the lower and upper probability bounds of the recognizing process. For instance, PST atom  $loc(id_1, c, 9)[.9, 1]$  says that a profile having  $id$   $id_1$  was in region  $c$  at time 9 with probability in the

interval  $[\cdot 9, 1]$  (the high-accuracy sensors used at security check points located in region  $c$  entail a narrow probability interval with upper bound equal to 1). Atom  $loc(id_1, a, 1)[\cdot 4, \cdot 7]$  says that  $id_1$  was recognized in region  $a$  at the earlier time 1 with probability in  $[\cdot 4, \cdot 7]$ . Assume that the information is represented as the set of atoms in Figure 1(b), which includes the two atoms above.  $\square$

PST atoms can be used to represent the output of a process aimed at tracking objects on the basis of sensor measurements. Generally, sensors are characterized by a *likelihood* function providing the conditional probability of obtaining a measurement given the value of a parameter  $\theta$ , such as the distance between the tracked object and the sensor. For instance, the likelihood function  $l(\theta)$  can represent the probability of detecting an object that has a distance of  $\theta$  meters from the sensor position. However, a likelihood function is generally not a probability distribution when viewed as a function of  $\theta$ . We may have  $l(\theta_1) = \cdot 9$ ,  $l(\theta_2) = \cdot 4$ ,  $l(\theta_3) = \cdot 1$ , and  $l(\theta \geq \theta_3) = 0$ , with distances  $\theta_1 < \theta_2 < \theta_3$ . This information can be encoded by using PST atoms  $loc(id, r_i, t)[l(\theta_i), l(\theta_i)]$  for each region  $r_i$  determined by the distance  $\theta_i$  (more general probability intervals can be used if the likelihood values are not known exactly).<sup>2</sup> However, several object tracking techniques combine the information of the likelihood function with the prior position distribution to obtain a probability distribution over *Space*. PST atoms can represent this kind of information by defining a PST atom with a single probability for each point in *Space*.<sup>3</sup> We refer the reader to the Related Work section for a detailed discussion on object tracking techniques and the relationship with our framework.

Although PST atoms express much useful information, they cannot express additional knowledge that integrity constraints can provide. In this paper we add integrity constraints to the original PST framework to form PST KBs. The integrity constraints have the form of *spatio-temporal denial* formulas (*std* formulas for short). We will soon see that such formulas are expressive enough to capture a large set of conditions. Basically, an *std* formula is a universally quantified negation of conjunctions of st-atoms and built-in predicates. We note that *std* formulas are related to a subclass of the first-order formulas introduced by Doder, Grant, and Ognjanović (2013), except that they (the *std* formulas) allow built-in predicates as well. In any case, the focus of Doder et al. is the axiomatization of various probabilistic spatio-temporal logics.

**Definition 2** (Std-formula). *An std-formula is an expression of the form*

$$\forall \mathbf{X}, \mathbf{Y}, \mathbf{Z} \neg \left[ \left( \bigwedge_{i=1}^k loc(X_i, Y_i, Z_i) \right) \wedge \alpha(\mathbf{X}) \wedge \beta(\mathbf{Y}) \wedge \gamma(\mathbf{Z}) \right]$$

where:

- $\mathbf{X}$  is a set of object variables,  $\mathbf{Y}$  is a set of space variables, and  $\mathbf{Z}$  is a set of time variables;
- $loc(X_i, Y_i, Z_i)$ , with  $i \in [1..k]$ , are st-atoms, where the  $X_i, Y_i, Z_i$  may be variables or constants of the appropriate type, such that, if  $X_i$  (resp.,  $Y_i, Z_i$ ) is a variable, then it occurs in  $\mathbf{X}$  (resp.,  $\mathbf{Y}, \mathbf{Z}$ ). Moreover, each variable in  $\mathbf{X}, \mathbf{Y}$ , and  $\mathbf{Z}$  occurs in at least one st-atom  $loc(X_i, Y_i, Z_i)$ , with  $i \in [1..k]$ ;

---

2. We note that PST KBs resulting from PST atoms encoding information provided by the likelihood function may be inconsistent because of the fact that the likelihood function need not be a probability distribution. This will turn out to be clearer after introducing the formal semantics of PST KBs in Section 2.2.

3. PST KBs resulting from PST atoms encoding the output of tracking turn out to be consistent.

- $\alpha(\mathbf{X})$  is a conjunction of built-in predicates of the form  $X_i \diamond X_j$ , where  $X_i$  and  $X_j$  are either variables occurring in  $\mathbf{X}$  or ids in  $ID$ , and  $\diamond$  is an operator in  $\{=, \neq\}$ ;
- $\beta(\mathbf{Y})$  is a conjunction of built-in predicates of the form  $Y_i \diamond Y_j$ , where  $Y_i$  and  $Y_j$  are either variables occurring in  $\mathbf{Y}$  or regions (i.e., non-empty subsets of  $Space$ ), and  $\diamond$  is a comparison operator in  $\{=, \neq, ov, nov\}$  (where *ov* stands for “overlaps” and *nov* stands for “does not overlap”);
- $\gamma(\mathbf{Z})$  is a conjunction of built-in predicates of the form  $Z_i \diamond Z_j$  where each  $Z_i$  and  $Z_j$  is either a time value in  $T$  or a variable in  $\mathbf{Z}$  that may be followed by  $+n$  where  $n$  is a positive integer and  $\diamond$  is an operator in  $\{=, \neq, <, \geq\}$ .

**Example 2.** In our running example, in region  $c$  security checks on one individual at a time are performed. The constraint “there cannot be two distinct objects in region  $c$  at any time between 1 and 20” can be expressed by the following std-formula:

$$f_1 = \forall X_1, X_2, Z_1 \neg[loc(X_1, c, Z_1) \wedge loc(X_2, c, Z_1) \wedge X_1 \neq X_2 \wedge Z_1 \geq 1 \wedge 20 \geq Z_1].$$

Due to the distance and the several obstacles between the entrance and the exit, we also have the constraint “no object can reach region  $e$  starting from region  $a$  in less than 10 time units”, that can be expressed as:

$$f_2 = \forall X_1, Z_1, Z_2 \neg[loc(X_1, a, Z_1) \wedge loc(X_1, e, Z_2) \wedge Z_1 < Z_2 \wedge Z_2 < Z_1 + 10].$$

Moreover, as the security check on each individual takes at least 2 time units, we know that “object  $id$  can go away from region  $c$  only if it stayed there for at least 2 time units”, that can be expressed as:

$$f_3 = \forall Y_1, Y_2, Z_1, Z_2, Z_3 \neg[loc(id, Y_1, Z_1) \wedge loc(id, c, Z_2) \wedge loc(id, Y_2, Z_3) \wedge Y_1 nov c \wedge Y_2 nov c \wedge Z_2 = Z_1 + 1 \wedge Z_2 < Z_3 \wedge Z_2 + 2 \geq Z_3]. \quad \square$$

In our work later it will be useful to distinguish std-formulas based on the number ( $k$ ) of st-atoms in them. In particular, *unary* std-formulas have  $k = 1$  and *binary* std-formulas have  $k = 2$ . In Example 2 above  $f_1$  and  $f_2$  are *binary* std-formulas while  $f_3$  is a *ternary* std-formula.

In the initial SPOT framework (Parker et al., 2007a) only PST atoms were considered. Moreover, it was assumed that all points in  $Space$  are reachable from all other points by all objects. To overcome this limitation, Grant et al. (2010) extended the SPOT framework by introducing *reachability definitions*. A *reachability atom* is written as  $reachable_{id}(p, q)$  where  $id \in ID$  is an object id, and  $p, q \in Space$ . Intuitively, the reachability atom says that it is possible for the object  $id$  to reach location  $q$  from location  $p$  in one unit of time. Hence, what is reachable in one time unit depends not only on the locations  $p$  and  $q$ , but also the object  $id$ . As we now show, reachability can be expressed in our formalism as an integrity constraint. However, in order to formulate reachability in our framework of denial formulas, we need to deal with what is not reachable, rather than what is reachable.

**Example 3.** Let  $r$  be the region consisting of all points  $q$  that are not reachable from  $p$  in one time unit. The corresponding std-formula is:

$$\forall X_1, Z_1, Z_2 \neg[loc(X_1, \{p\}, Z_1) \wedge loc(X_1, r, Z_2) \wedge Z_2 = Z_1 + 1]. \quad \square$$

This integrity constraint is used only if not all points are reachable from  $p$  in one time unit. We can also express which points can not be reached from  $p$  in any number of time units, not just 1, by changing  $Z_1 + 1$  to  $Z_1 + i$ .

**Example 4.** In our running example, the following std-formula states that the points in region  $r = \{(x, y) | 0 \leq x \leq 5 \wedge y = 3\}$  (i.e., those close to the upper-side of the wall dividing the hall room and the one where there are queue dividers) are not reachable in less than 3 time units from any point in  $r' = \{(x, y) | 0 \leq x \leq 5 \wedge y = 2\}$  (i.e., the points close to the other side of that wall):  $f_4 = \forall X_1, Z_1, Z_2 \neg [loc(X_1, r', Z_1) \wedge loc(X_1, r, Z_2) \wedge Z_1 < Z_2 \wedge Z_2 < Z_1 + 3]$ .  $\square$

We are now ready to formally define PST KBs.

**Definition 3 (PST KB).** Given the sets  $ID$ ,  $T$ , and  $Space$ , a PST KB  $\mathcal{K}$  is a pair  $\langle \mathcal{A}, \mathcal{F} \rangle$ , where  $\mathcal{A}$  is a finite set of PST atoms and  $\mathcal{F}$  is finite set of std-formulas using object symbols in  $ID$ , time values in  $T$ , and spatial regions consisting of sets of points in  $Space$ .

**Example 5.** In our running example,  $ID = \{id_1, id_2\}$ ,  $T = [0, 20]$ ,  $Space$  is the set of points  $(x, y)$  such that  $0 \leq x \leq 7$  and  $0 \leq y \leq 7$ , and PST KB  $\mathcal{K}_{ex}$  is the pair  $\langle \mathcal{A}_{ex}, \mathcal{F}_{ex} \rangle$ , where  $\mathcal{A}_{ex}$  is the set consisting of the PST atoms in Figure 1(b), and  $\mathcal{F}_{ex}$  is the set  $\{f_1, f_2, f_3, f_4\}$  of std-formulas defined in Examples 2 and 4.  $\square$

## 2.2 Semantics

The semantics of a PST KB is defined through the concept of *worlds*. Before introducing this concept, we define *ground* std-formulas.

Given an std-formula  $f$  having the form in Definition 2, we denote by  $\Theta_f$  the set of all substitutions of variables in  $\mathbf{X}$ ,  $\mathbf{Y}$ , and  $\mathbf{Z}$  with constants in  $ID$ ,  $\mathcal{S}$ , and  $T$ , respectively, where  $\mathcal{S}$  is the set of all subsets of  $Space$  that contain a single point.<sup>4</sup> Moreover, given substitution  $\theta \in \Theta_f$ , we denote as  $\theta(f)$  the *ground* std-formula resulting from applying  $\theta$  to  $f$ :  $\theta(f) = \neg[(\bigwedge_{i=1}^k loc(\theta(X_i), \theta(Y_i), \theta(Z_i))) \wedge \alpha(\theta(\mathbf{X})) \wedge \beta(\theta(\mathbf{Y})) \wedge \gamma(\theta(\mathbf{Z}))]$ . The ground conjunction of built-in predicates  $\alpha(\theta(\mathbf{X})) \wedge \beta(\theta(\mathbf{Y})) \wedge \gamma(\theta(\mathbf{Z}))$  evaluates to either *true* or *false*. When it is true we omit it. So  $\theta(f)$  is either the negation of a conjunction of ground st-atoms or the truth value *true* (when the conjunction of built-in predicates evaluates to *false*).

**Example 6.** Consider the formula  $f_1 = \forall X_1, X_2, Z_1 \neg [loc(X_1, c, Z_1) \wedge loc(X_2, c, Z_1) \wedge X_1 \neq X_2 \wedge Z_1 \geq 1 \wedge 20 \geq Z_1]$  introduced in Example 2, and the substitution  $\theta = \{X_1/id_1, X_2/id_2, Z_1/6\}$ , where  $id_1, id_2 \in ID$  and  $6 \leq tmax$ . Thus,  $\theta(f_1) = \neg [loc(id_1, c, 6) \wedge loc(id_2, c, 6)]$ , where the conjunction of ground built-in predicates  $id_1 \neq id_2 \wedge 6 \geq 1 \wedge 6 \leq 20$ , evaluating to *true*, is not reported in  $\theta(f_1)$ .  $\square$

**Definition 4 (World).** A world  $w$  is a function,  $w : ID \times T \rightarrow Space$ .

Basically, a world  $w$  specifies a trajectory for each  $id \in ID$ . That is, for each  $id \in ID$ ,  $w$  says where in  $Space$  object  $id$  was/is/will be at each time  $t \in T$ . In particular, this means that an object can be in only one location at a time.<sup>5</sup> However, a location may contain multiple objects. It is easy to see that world  $w$  can be represented by the set  $\{loc(id, \{p\}, t) | w(id, t) = p\}$  of ground st-atoms.

4. We use only such singleton subsets of  $Space$  in order to reduce the number of possible instantiations of variables  $Y$  from exponential to linear in the size of  $Space$ , without serious effect on the meanings of the std-formulas.

5. In some examples it may be useful to allow objects to enter or leave the space under consideration. This can be accomplished, for instance, by having one or more external points outside of the space where objects may be located. To simplify matters we assume that  $Space$  contains all these points.

**Example 7.** World  $w_1$  describing the trajectories of  $id_1$  and  $id_2$  for time units in  $[0, 20]$  is such that  $w_1(id_1, t) = (4, 1)$  for  $t \in [0, 5]$ ,  $w_1(id_1, t) = (7, 2)$  for  $t \in [6, 7]$ ,  $w_1(id_1, t) = (7, 4)$  for  $t \in [8, 10]$ ,  $w_1(id_1, t) = (4, 4)$  for  $t \in [11, 16]$ ,  $w_1(id_1, t) = (1, 6)$  for  $t \in [17, 20]$ ,  $w_1(id_2, t) = (4, 1)$  for  $t \in [0, 11]$ ,  $w_1(id_2, t) = (7, 5)$  for  $t \in [12, 15]$ ,  $w_1(id_2, t) = (7, 7)$  for  $t \in [16, 16]$ ,  $w_1(id_2, t) = (4, 5)$  for  $t \in [17, 20]$ .  $\square$

**Definition 5 (Satisfaction).** Given a world  $w$  and a ground st-atom  $a = loc(id, r, t)$ , we say that  $w$  satisfies  $a$  (denoted as  $w \models a$ ) iff  $w(id, t) \in r$ . Moreover, we say that  $w$  satisfies a conjunction of ground st-atoms  $\bigwedge_{i=1}^k a_i$  (denoted as  $w \models \bigwedge_{i=1}^k a_i$ ) iff  $w \models a_i \forall i \in [1..k]$ . Finally, world  $w$  satisfies std-formula  $f$  (denoted as  $w \models f$ ) iff for each substitution  $\theta \in \Theta_f$ ,  $w \models \theta(f)$ .

Note that, as there is a negation in front of  $f$ ,  $w \models \theta(f)$  iff  $w$  does not satisfy a ground st-atom in  $\theta(f)$  or the conjunction of ground built-in predicates in  $\theta(f)$  evaluates to false.

**Example 8.** World  $w_1$  of Example 7 satisfies the st-atom  $loc(id_1, b, 0)$ , as  $w_1(id_1, 0) = (4, 1)$  belongs to region  $b$  (see Figure 1(a)). Moreover,  $w_1 \models \neg[loc(id_1, b, 0) \wedge loc(id_1, e, 15)]$  as  $w_1 \not\models loc(id_1, e, 15)$ , since  $w_1(id_1, 15) = (4, 4) \notin e$ .  $\square$

In the following, we will denote as  $\mathcal{W}(\mathcal{K})$  the set of all worlds of the PST KB  $\mathcal{K}$ . Moreover, in order to simplify formulas, we will assume that  $w$  ranges over  $\mathcal{W}(\mathcal{K})$ .

An interpretation  $I$  for a PST KB  $\mathcal{K}$  is a probability distribution function (PDF) over  $\mathcal{W}(\mathcal{K})$ , that is, a function assigning a probability value to each world in  $\mathcal{W}(\mathcal{K})$ .  $I(w)$  is the probability that  $w$  describes the actual trajectories of all the objects.<sup>6</sup> Some interpretations are models of  $\mathcal{K}$  in which case we write  $M$  instead of  $I$ .

**Definition 6 (Model).** A model  $M$  for a PST KB  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  is an interpretation for  $\mathcal{K}$  such that:

- $\forall loc(id, r, t)[\ell, u] \in \mathcal{A}, \sum_{w \mid w \models loc(id, r, t)} M(w) \in [\ell, u];$
- $\forall f \in \mathcal{F}, \sum_{w \mid w \not\models f} M(w) = 0.$

The first condition in the definition above means that, for each atom  $a = loc(id, r, t)[\ell, u] \in \mathcal{A}$ , the sum of the probabilities assigned by  $M$  to the worlds satisfying the st-atom  $loc(id, r, t)$  have to belong to the probability interval  $[\ell, u]$  specified by  $a$ . The second condition means that every world not satisfying a formula  $f \in \mathcal{F}$  must be assigned by  $M$  probability equal to 0.

**Example 9.** Let  $w_1$  be the world introduced in Example 7. Let  $w_2$  be as  $w_1$  except that  $w_2(id_1, 1) = (3, 2)$ , and let  $w_3$  be as  $w_2$  except that  $w_3(id_2, 2) = (2, 2)$ ,  $w_3(id_2, t) = (0, 3)$  for  $t \in [18..20]$ . Let  $M$  be such that  $M(w_1) = .7$   $M(w_2) = .2$   $M(w_3) = .1$ , and  $M(w) = 0$  for all the other worlds in  $\mathcal{W}(\mathcal{K}_{ex})$ . It can be checked that  $M$  satisfies both conditions of Definition 6 for the PST KB  $\mathcal{K}_{ex}$  of our running example. For instance, for atom  $loc(id_1, a, 1)[.4, .7] \in A_{ex}$ ,  $\sum_{w \mid w \models loc(id_1, a, 1)} M(w) = M(w_1) = .7 \in [.4, .7]$  (note that, at time 1,  $w_2(id_1, 1) = w_3(id_1, 1) = (3, 2)$  that is not in region  $a$ ). Moreover, it is easy to check that  $w_1, w_2, w_3$  satisfy every std-formula in  $\mathcal{F}_{ex}$ . Thus,  $M$  is a model for  $\mathcal{K}_{ex}$ .  $\square$

We say that PST KB  $\mathcal{K}$  is consistent iff if there is a model for it. The set of models for  $\mathcal{K}$  will be denoted as  $\mathbf{M}(\mathcal{K})$ .

6. As a PDF,  $I(w)$  is non-negative and sums to 1 over worlds.

**Definition 7** (Consistency). PST KB  $\mathcal{K}$  is consistent iff  $\mathbf{M}(\mathcal{K}) \neq \emptyset$ .

**Example 10.** PST KB  $\mathcal{K}_{ex}$  of our running example is consistent, as there exists the model  $M$  of Example 9 for it.  $\square$

### 3. Checking the Consistency of PST KBs

In this section, we address the fundamental problem of checking the consistency of PST KBs. Given a PST KB  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$ , the consistency checking problem is deciding whether  $\mathbf{M}(\mathcal{K}) \neq \emptyset$ , that is, whether there is a model for  $\mathcal{K}$ .

In Section 3.1 we show that the consistency checking problem is NP-complete. Our goal in the rest of this section is to find more efficient ways to determine consistency. In Section 3.2 we find a sufficient condition using a set of mixed-binary linear inequalities. Then, in Section 3.3 we find a necessary condition using a different set of linear inequalities. We deal with the special case where all the std-formulas are unary in Section 3.4. Finally we investigate in detail the case where the std-formulas are either unary or binary in Section 3.5.

For our complexity analysis we take the size of a PST KB  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$ , whose PST atoms and std-formulas are built from the constants in  $ID$ ,  $T$ , and  $Space$ , as the number of PST atoms and std-formulas in  $\mathcal{K}$  plus the number of items in  $ID$ ,  $T$ , and  $Space$ , that is,  $|\mathcal{K}| = |\mathcal{A}| + |\mathcal{F}| + |ID| + |T| + |Space|$ .

#### 3.1 Checking Consistency is NP-Complete

Before considering the case of general PST KBs, we first note that the consistency checking problem was addressed in the initial SPOT framework of Parker et al. (2007a) where only PST atoms were considered. This is the special case of our PST KB concept where  $\mathcal{F} = \emptyset$ . It was shown there that the consistency of a PST KB  $\mathcal{K} = \langle \mathcal{A}, \emptyset \rangle$  (using our notation) can be checked in polynomial time w.r.t. the size of  $\mathcal{K}$  by solving a set of linear inequalities whose variables  $v_{id,t,p}$  represent the probability that object  $id$  is at point  $p$  at time  $t$ .

The reason for presenting the result of Parker et al. is twofold: first, we compare the complexity of the consistency checking problem for general PST KBs with that of the initial SPOT framework; second, we use it to prove some tractability results for PST KBs.

**Fact 1** (Parker et al.). Let  $\mathcal{K} = \langle \mathcal{A}, \emptyset \rangle$  be a PST KB (where the set of std-formulas is empty). Then  $\mathcal{K}$  is consistent iff there is a feasible solution of  $CC(\mathcal{K})$ , where  $CC(\mathcal{K})$  consists of the following (in)equalities:

- (1)  $\forall loc(id, r, t)[\ell, u] \in \mathcal{A}: \ell \leq \sum_{p \in r} v_{id,t,p} \leq u;$
- (2)  $\forall id \in ID, t \in T: \sum_{p \in Space} v_{id,t,p} = 1;$
- (3)  $\forall id \in ID, t \in T, p \in Space: v_{id,t,p} \geq 0.$

Basically, inequalities (1) ensure that a solution of  $CC(\mathcal{K})$  places the object in  $r$  with a probability between  $\ell$  and  $u$ , as required by the atom  $(id, r, t, [\ell, u])$ . Inequalities (2) and (3) ensure that for each  $id$  and  $t$ , the  $v_{id,t,p}$  variables jointly represent a probability distribution. Fact 1 is correct because every model  $M$  for  $\mathcal{K}$  corresponds to a solution  $\sigma$  for  $CC(\mathcal{K})$  such that the sum of the

probabilities assigned by  $M$  to the worlds for  $\mathcal{K}$  satisfying an st-atom  $loc(id, \{p\}, t)$  is equal to the value assigned to variable  $v_{id,t,p}$  by  $\sigma$ .

Now we state our first result: the consistency checking problem is NP-complete.

**Theorem 1.** *Given a PST KB  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$ , deciding whether  $\mathcal{K}$  is consistent is NP-complete.*

*Proof. (Membership).* We show that checking the consistency of  $\mathcal{K}$  can be reduced to deciding an instance  $\phi_{\mathcal{K}}$  of (an extension to) the Probabilistic Satisfiability (PSAT) problem (Hailperin, 1984; Nilsson, 1986), which is in NP (Georgakopoulos, Kavvadias, & Papadimitriou, 1988). Given a set of  $m$  clauses  $C_1, \dots, C_m$ , each of them consisting of a disjunction of one or more literals constructed from the propositional variables  $x_1, \dots, x_n$ , and probability values  $\Pr(C_1), \dots, \Pr(C_m)$  for each clause, PSAT is the problem of deciding whether there is a probability distribution  $\omega$  over the set of the  $2^n$  truth assignments for the propositional variables  $x_1, \dots, x_n$  such that for each clause  $C_i$ , the sum of the probabilities assigned by  $\omega$  to the truth assignments satisfying  $C_i$  is equal to  $\Pr(C_i)$ , with  $i \in [1..m]$ . PSAT is a generalization of SAT, which is obtained from PSAT by assigning probability equal to one to each clause. Georgakopoulos et al. formulated PSAT in terms of the feasibility of a system of  $m + 1$  linear equations using  $2^n$  variables corresponding to the probabilities assigned by  $\omega$  to truth assignments. To show the existence of polynomial-size witness, the following result from linear programming theory was exploited by Georgakopoulos et al.: if a system of  $\kappa$  linear equalities has a feasible solution, then it admits at least one feasible solution with at most  $\kappa$  non-zero variables (Papadimitriou & Steiglitz, 1982). In what follows we consider an extension to PSAT where each clause  $C_i$  is associated with a probability interval  $[\Pr_{\ell}(C_i), \Pr_u(C_i)]$ , instead of a single value. The membership in NP for this extension straightforwardly follows from the membership proof provided by Georgakopoulos et al. for PSAT, as using probability intervals can be still formulated as the linear system introduced by Georgakopoulos et al. after reducing double-sided inequalities to equalities with single bounded slack variables (Jaumard, Hansen, & de Aragão, 1991).

Given a PST KB  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$ , we define an instance  $\phi_{\mathcal{K}}$  of PSAT where each clause is associated with a probability interval. Let  $U$  be the set of all propositional variables  $x_{id,p,t}$  such that  $id \in ID$ ,  $p \in Space$ , and  $t \in T$  (i.e., each st-atom  $loc(id, \{p\}, t)$  corresponds to propositional variable  $x_{id,p,t} \in U$ ). The conjunction  $\phi_{\mathcal{K}}$  of clauses associated with probability intervals is defined as follows:

- For each PST atom  $a = loc(id, r, t)[\ell, u] \in \mathcal{A}$ ,  $\phi_{\mathcal{K}}$  consists of the clause  $C_a = \bigvee_{p \in r} x_{id,p,t}$ . The probability interval  $[\Pr_{\ell}(C_a), \Pr_u(C_a)]$  is equal to  $[\ell, u]$ .
- For each  $f \in \mathcal{F}$  and  $\theta \in \Theta_f$  such that  $\theta(f) = \neg[\bigwedge_{i=1}^k loc(\theta(X_i), \theta(Y_i), \theta(Z_i))]$ ,  $\phi_{\mathcal{K}}$  consists of clause  $C_{\theta(f)} = \bigvee_{i=1}^k \neg x_{\theta(X_i), \theta(Y_i), \theta(Z_i)}$ , whose probability interval is  $[1, 1]$ .
- For each  $id \in ID$  and  $t \in T$ ,  $\phi_{\mathcal{K}}$  consists of the clause  $C_{id,t} = \bigvee_{p \in Space} x_{id,p,t}$ , whose probability interval is  $[1, 1]$ .
- For each  $id \in ID$ ,  $t \in T$ , and  $p_i, p_j \in Space$ , with  $p_i \neq p_j$ ,  $\phi_{\mathcal{K}}$  consists of the clause  $C_{id,t,p_i,p_j} = \neg x_{id,p_i,t} \vee \neg x_{id,p_j,t}$  whose probability interval is  $[1, 1]$ .

It is easy to see that for each world  $w \in \mathcal{W}(\mathcal{K})$ , there is a truth assignment  $\tau_w$  for the variables in  $U$  such that  $\tau_w(x_{id,p,t})$  is *true* iff  $w(id, t) = p$ . However, there are truth assignments which do not correspond to any world in  $\mathcal{W}(\mathcal{K})$  (for instance, those where both  $\tau(x_{id,p_i,t})$  and  $\tau(x_{id,p_j,t})$ , with  $p_i \neq p_j$ , are *true*). We now show that  $\mathcal{K}$  is consistent iff  $\phi_{\mathcal{K}}$  is satisfiable.

( $\Rightarrow$ ) Given a model  $M$  for  $\mathcal{K}$ , we show that there is a PDF  $\omega$  over the set of truth assignments such that for each clause  $C$  of  $\phi_{\mathcal{K}}$ , the sum of the probabilities assigned by  $\omega$  to the truth assignments satisfying  $C$  belongs to  $[\text{Pr}_\ell(C), \text{Pr}_u(C)]$ . Let  $\omega$  be such that for each truth assignment  $\tau_w$  corresponding to a world  $w \in \mathcal{W}(\mathcal{K})$ ,  $\omega(\tau_w) = M(w)$  and  $\omega(\tau) = 0$  for all other truth assignments  $\tau$  not corresponding to a world. It is easy to check that the conditions of Definition 6 entail that clauses of the form  $C_a$  and  $C_{\theta(f)}$  are satisfied by  $\omega$ , and clauses of the form  $C_{id,t}$  and  $C_{id,t,p_i,p_j}$  are satisfied as well since only  $\tau_{w_1}, \dots, \tau_{w_{|\mathcal{W}(\mathcal{K})}}$  may have been assigned by  $\omega$  a probability different from 0 and every world  $w \in \mathcal{W}(\mathcal{K})$  by definition assigns exactly one point in  $Space$  to each  $id, t$  pair.

( $\Leftarrow$ ) Let  $\omega$  be a PDF over the set of truth assignments such that for each clause  $C$  of  $\phi_{\mathcal{K}}$ , the sum of probabilities assigned by  $\omega$  to the truth assignments satisfying  $C$  belongs to  $[\text{Pr}_\ell(C), \text{Pr}_u(C)]$ . A model  $M$  for  $\mathcal{K}$  can be defined as  $M(w) = \omega(\tau_w)$  where  $\tau$  is the truth assignment corresponding to world  $w \in \mathcal{W}(\mathcal{K})$ . Since the clauses of the form  $C_{id,t}$  (resp.,  $C_{id,t,p_i,p_j}$ ) are satisfied by  $\omega$ , all truth assignments  $\tau$  such that  $\tau(x_{id,p,t}) = \text{false}$  for all  $p \in Space$  (resp.,  $\tau(x_{id,p_i,t}) = \text{true}$  and  $\tau(x_{id,p_j,t}) = \text{true}$ , with  $p_i, p_j \in Space$ ,  $p_i \neq p_j$ ) are assigned by  $\omega$  probability 0. Hence, all other truth assignments correspond to a world, and the fact that the clauses of the form  $C_a$  and  $C_{\theta(f)}$  are satisfied by  $\omega$  entails that the conditions of Definition 6 hold.

(*Hardness*). We show a reduction to our problem from the NP-hard Hamiltonian path problem (Papadimitriou, 1994), that is, the problem of checking whether there is a path  $\pi$  in a directed graph  $G$  such that  $\pi$  visits each vertex of  $G$  exactly once.

Given a directed graph  $G = \langle V, E \rangle$ , where  $V = \{v_0, \dots, v_k\}$  is the set of its vertices, and  $E$  is a set of pairs  $(v_i, v_j)$  with  $v_i, v_j \in V$ , we construct an instance of our problem as follows. Let  $ID = \{id\}$ ,  $Space = V$ , and  $T = [0, \dots, k]$ .  $\mathcal{K}$  is the pair  $\langle \mathcal{A}, \mathcal{F} \rangle$  such that  $\mathcal{A}$  consists of the PST atom  $loc(id, \{v_0\}, 0)[1, 1]$  and  $\mathcal{F}$  consists of std-formulas  $f_1^i$  (with  $i \in [0..k]$ ) and  $f_2$  such that:

- $f_1^i = \forall Z_1, Z_2 \neg [loc(id, \{v_i\}, Z_1) \wedge loc(id, Space \setminus V', Z_2) \wedge Z_2 = Z_1 + 1]$  where  $V'$  is the set of vertices  $v_j$  s.t.  $(v_i, v_j) \in E$ . This formula says that the only points  $id$  can reach starting from  $v_i$  in one time step are those in  $V'$ . ( $f_1^i$  does not exist if  $V' = Space$ .)
- $f_2 = \forall Y_1, Z_1, Z_2 \neg [loc(id, Y_1, Z_1) \wedge loc(id, Y_1, Z_2) \wedge Z_1 \neq Z_2]$ , saying that  $id$  can not be at the same location for distinct time values.

We show that  $\mathcal{K}$  is consistent iff there is a Hamiltonian path in  $G$ .

( $\Rightarrow$ ) As there is only one  $id$  in  $\mathcal{A}$ , every world  $w \in \mathcal{W}(\mathcal{K})$  is such that  $w$  places  $id$  on a vertex in  $V$  at each time value  $t \in T$ . As  $\mathcal{K}$  is consistent, there is a model  $M \in \mathbf{M}(\mathcal{K})$  such that  $M$  assigns probability greater than zero only to worlds  $w$  such that  $\forall f \in \mathcal{F}, w \models f$ . In particular, let  $w$  be one such world. The fact that  $w \models f_1^i$  entails that  $\forall t \in [0, k-1], w(id, t) = v_i$  and  $w(id, t+1) = v_j$  iff  $(v_i, v_j) \in E$ . Moreover, the fact that  $w \models f_2$  entails that  $\forall t, t' \in [0, k], t \neq t', w(id, t) \neq w(id, t')$ , meaning that  $id$  is never placed by  $w$  on the same vertex at different time units. Since  $loc(id, v_0, 0)[1, 1] \in \mathcal{A}$ , every world which is assigned probability greater than zero by  $M$  is such that  $w(id, 0) = v_0$ . It follows that every world  $w \in \mathcal{W}(\mathcal{K})$  which is assigned by  $M \in \mathbf{M}(\mathcal{K})$  a probability greater than zero encodes a Hamiltonian path of  $G$  whose first vertex is  $v_0$ . In fact,  $\forall w \in \mathcal{W}(\mathcal{K})$  such that  $M(w) > 0$  the following properties hold: (i)  $w(id, 0) = v_0$ ,

(ii)  $\forall t \in [0, k-1]$ ,  $w(id, t) = v_i$ ,  $w(id, t+1) = v_j$  iff  $(v_i, v_j) \in E$ . (iii)  $\forall t, t' \in [0, k]$ ,  $t \neq t'$ ,  $w(id, t) \neq w(id, t')$ . Conditions (i) and (ii) entail that  $\pi = w(id, 0), w(id, 1), \dots, w(id, k)$  is a path on  $G$  starting from vertex  $v_0$ , while condition (iii) entails that each vertex  $v \in V$  occurs exactly once in  $\pi$ .

( $\Leftarrow$ ) Let  $\pi$  be a Hamiltonian path of  $G$ . We denote by  $\pi[i]$  (with  $i \in [0..k]$ ) the  $i$ -th vertex of  $\pi$ . W.l.o.g. we assume that the first vertex of  $\pi$  is  $v_0$ , that is,  $\pi[0] = v_0$ . We now show that  $\mathcal{K}$  is consistent by finding a model for it. Let  $M$  be a function over  $\mathcal{W}$  such that for all worlds  $w \in \mathcal{W}$ ,  $M(w) = 0$ , except for the world  $w^*$  which is such that:  $w^*(id, 0) = \pi[0] = v_0$ ,  $\forall t \in [1, k]$ ,  $w^*(id, t) = \pi[t]$ . It is easy to see that  $w^* \models \mathcal{F}$ . In fact, for each  $i \in [0..k]$ ,  $f_1^i$  is satisfied by  $w^*$ , since the fact that  $\pi$  is a path on  $G$  entails that  $\forall t \in [0, k-1]$ ,  $w^*(id, t) = v_i$  and  $w^*(id, t+1) = v_j$  only if edge  $(v_i, v_j)$  is an edge of  $G$ . Moreover,  $f_2$  is satisfied by  $w^*$ , since the fact that  $\pi$  is a Hamiltonian path entails that  $w^*$  places  $id$  on different locations (i.e., vertices of  $G$ ) at different times. Since  $w^* \models \mathcal{F}$ , it can be assigned by  $M$  a probability different from 0. Let  $M(w^*) = 1$ . Therefore, as  $\sum_{w|w \models loc(id, v_0, 0)} M(w) = M(w^*) + \sum_{w|w \neq w^* \wedge w \models loc(id, v_0, 0)} M(w) = 1$ , the condition required by  $atom\ loc(id, v_0, 0)[1, 1] \in \mathcal{A}$  holds too. Thus,  $M$  is a model for  $\mathcal{K}$ .  $\square$

We note that NP-hardness holds already for binary std-formulas. In Section 3.5 we will find conditions that make the consistency checking problem tractable for binary std-formulas. The reduction shown in the membership proof, from the consistency checking problem to PSAT, would allow us to define additional tractable cases if the PSAT instances resulting from such a reduction were tractable. However, as will be discussed in Section 6, the tractable cases that have been identified for PSAT (Georgakopoulos et al., 1988; Andersen & Pretolani, 2001) do not carry over to our framework.

### 3.2 Sufficient Condition for Checking Consistency

We present a set of mixed-binary linear inequalities whose feasibility entails the consistency of PST KB  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$ . As explained in Section 3.1, Parker et al. (2007a) showed that the consistency of a PST KB  $\mathcal{K} = \langle \mathcal{A}, \emptyset \rangle$  (using our notation) can be checked in polynomial time w.r.t. the size of  $\mathcal{K}$  by solving a set of linear inequalities whose variables  $v_{id,t,p}$  represent the probability that object  $id$  is at point  $p$  at time  $t$ . Here, we start from this set of linear inequalities and augment it with some inequalities ensuring that if the so-obtained set of linear inequalities has a feasible solution then every ground std-formula derived from  $\mathcal{F}$  is satisfied. To achieve this, we need to introduce the binary variables  $\delta_i$ , thus obtaining a set of mixed-binary linear inequalities.

**Definition 8** (MBL( $\mathcal{K}$ )). *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$ . MBL( $\mathcal{K}$ ) consists of the following (in)equalities:*

$$(1) \forall loc(id, r, t)[\ell, u] \in \mathcal{A}: \ell \leq \sum_{p \in r} v_{id,t,p} \leq u;$$

$$(2) \forall id \in ID, t \in T: \sum_{p \in Space} v_{id,t,p} = 1;$$

$$(3) \forall id \in ID, t \in T, p \in Space: v_{id,t,p} \geq 0;$$

$$(4) \text{ for each } f \in \mathcal{F} \text{ and } \theta \in \Theta_f \text{ such that } \theta(f) = \neg \left[ \bigwedge_{i=1}^k loc(\theta(X_i), \theta(Y_i), \theta(Z_i)) \right], \text{ where } k \text{ is the number of st-atoms in } f, \text{ the (in)equalities:}$$

- (a)  $\forall i \in [1..k] : \sum_{p \in \theta(Y_i)} v_{\theta(X_i), \theta(Z_i), p} \leq \delta_i;$   
 (b)  $\sum_{i=1}^k \delta_i = k - 1;$   
 (c)  $\forall i \in [1..k] : \delta_i \in \{0, 1\}.$

Basically, inequalities (1) ensure that a solution of  $MBL(\mathcal{K})$  places the object in  $r$  with a probability between  $\ell$  and  $u$ , as required by the atom  $(id, r, t, [\ell, u])$ . Inequalities (2) and (3) ensure that for each  $id$  and  $t$ , the  $v_{id,t,p}$  variables jointly represent a probability distribution. Moreover, for each ground st-atom  $loc(\theta(X_i), \theta(Y_i), \theta(Z_i))$  of the ground std-formula  $\theta(f)$ , inequalities (4)(a) and (4)(c) entail that the probability  $v_{\theta(X_i), \theta(Z_i), p}$  that object  $\theta(X_i)$  is in any point  $p$  in region  $\theta(Y_i)$  at time  $\theta(Z_i)$  is either constrained to be 0 or free to take any value not greater than 1. Intuitively enough, if  $v_{\theta(X_i), \theta(Z_i), p}$  is enforced to be zero (i.e.,  $\delta_i = 0$ ), then object  $\theta(X_i)$  can not be in region  $\theta(Y_i)$  at time  $\theta(Z_i)$ . On the other hand, if  $v_{\theta(X_i), \theta(Z_i), p}$  is left free to take any value less than or equal to one (i.e.,  $\delta_i = 1$ ), then  $\theta(X_i)$  may or may not be in region  $\theta(Y_i)$  at time  $\theta(Z_i)$ . Finally, equality (4)(b) entails that there is at least one of the  $k$  ground st-atoms  $loc(\theta(X_i), \theta(Y_i), \theta(Z_i))$  of  $\theta(f)$  such that  $\theta(X_i)$  is not placed in a point in  $\theta(Y_i)$  at time  $\theta(Z_i)$ .

**Example 11.** Consider the ground std-formula  $\theta(f_1) = \neg[loc(id_1, c, 6) \wedge loc(id_2, c, 6)]$  of Example 6. Then, the inequalities in  $MBL(\mathcal{K})$  corresponding to  $\theta(f_1)$  are:

$$(4a) \sum_{p \in c} v_{id_1, 6, p} \leq \delta_1; \quad \sum_{p \in c} v_{id_2, 6, p} \leq \delta_2; \quad (4b) \delta_1 + \delta_2 = 1; \quad (4c) \delta_1, \delta_2 \in \{0, 1\}. \quad \square$$

The following theorem states that  $MBL(\mathcal{K})$  can be used to check if  $\mathcal{K}$  is consistent.

**Theorem 2.** *If there is a feasible solution of  $MBL(\mathcal{K})$  then  $\mathcal{K}$  is consistent.*

*Proof.* Let  $\sigma$  be a solution of  $MBL(\mathcal{K})$ , and  $\sigma(v_{id,t,p})$  the value assigned to variable  $v_{id,t,p}$  by  $\sigma$ . We define the function  $M$  over  $\mathcal{W}(\mathcal{K})$  such that, for each world  $w \in \mathcal{W}(\mathcal{K})$ ,  $M(w) = \prod_{id \in ID, t \in T, w(id,t)=p} \sigma(v_{id,t,p})$ , that is  $M(w)$  is the product of the values assigned by solution  $\sigma$  to variables  $v_{id,t,p}$  such that  $w(id, t) = p$ . It can be shown that, (in)equalities (2) and (3) of the definition of  $MBL(\mathcal{K})$  entail that  $M$  is a PDF over  $\mathcal{W}(\mathcal{K})$ . Moreover, since  $\sigma(v_{id,t,p})$  is equal to  $\sum_{w|w \models loc(id,t,p)} M(w)$ , for each atom  $loc(id, r, t)[\ell, u] \in \mathcal{A}$ ,  $\sum_{w|w \models loc(id,r,t)} M(w) = \sum_{p \in r} \sum_{w|w \models loc(id,t,p)} M(w) = \sum_{p \in r} \sigma(v_{id,t,p}) \in [\ell, u]$ . Given  $f \in \mathcal{F}$  and  $\theta \in \Theta_f$  such that  $\theta(f)$  is logically equivalent to the negation of the conjunction of the st-atoms  $\bigwedge_{i=1}^k loc(\theta(X_i), \theta(Y_i), \theta(Z_i))$ , the inequalities (4)(a-c) entail that there is  $i \in [1..k]$  such that  $\sum_{p \in \theta(Y_i)} \sigma(v_{\theta(X_i), \theta(Z_i), p}) = 0$ . Thus  $\forall p \in \theta(Y_i), \sigma(v_{\theta(X_i), \theta(Z_i), p}) = 0$ . Hence, for each world  $w \in \mathcal{W}(\mathcal{K})$  such that  $w(\theta(X_i), \theta(Z_i)) = p$ ,  $M(w) = 0$  due to the presence of the factor  $\sigma(v_{\theta(X_i), \theta(Z_i), p}) = 0$  in the product defining  $M(w)$ . Therefore, for each std-formula  $f \in \mathcal{F}$ ,  $\sum_{w|w \not\models f} M(w) = 0$ ; hence  $M$  is a model for  $\mathcal{K}$  and so  $\mathcal{K}$  is consistent.  $\square$

A consequence of Theorem 2 is that well-known techniques for solving linear optimization problems can be adopted to address the consistency checking problem, thus taking advantage of results from more than 50 years of research on integer linear programming (Jünger et al., 2010).

The following example shows that the converse of Theorem 2 does not hold ( $\mathcal{K}$  may be consistent even if there is no feasible solution of  $MBL(\mathcal{K})$ ).

**Example 12.** Let  $ID = \{id\}$ ,  $T = [0, 1]$ ,  $Space = \{p_0, p_1\}$ , and  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  where:  
 $\mathcal{A} = \{loc(id, \{p_0\}, 0)[0.5, 0.5], loc(id, \{p_1\}, 1)[0.5, 0.5]\}$  and  
 $\mathcal{F} = \{\neg[loc(id, \{p_0\}, 0) \wedge loc(id, \{p_1\}, 1)]\}.$

Thus,  $\mathcal{W} = \{w_1, w_2, w_3, w_4\}$  where:  $w_1(id, 0) = p_0, w_1(id, 1) = p_0, w_2(id, 0) = p_0, w_2(id, 1) = p_1, w_3(id, 0) = p_1, w_3(id, 1) = p_0, w_4(id, 0) = p_1, w_4(id, 1) = p_1$ .

It is easy to check that the  $M$  for which  $M(w_1) = 0.5, M(w_2) = 0, M(w_3) = 0, M(w_4) = 0.5$  is a model for  $\mathcal{K}$ . Then  $MBL(\mathcal{K})$  includes the following inequalities:

$$0.5 \leq v_{id,0,p_0} \leq 0.5; \quad 0.5 \leq v_{id,1,p_1} \leq 0.5; \quad v_{id,0,p_0} + v_{id,0,p_1} = 1; \quad v_{id,1,p_0} + v_{id,1,p_1} = 1;$$

$$v_{id,0,p_0} \leq \delta_1; \quad v_{id,1,p_1} \leq \delta_2; \quad \delta_1 + \delta_2 = 1; \quad \delta_1, \delta_2 \in \{0, 1\};$$

$$v_{id,0,p_0} \geq 0, v_{id,0,p_1} \geq 0, v_{id,1,p_0} \geq 0, v_{id,1,p_1} \geq 0.$$

The first two inequalities force both  $v_{id,0,p_0}$  and  $v_{id,1,p_1}$  to be 0.5. So in the second line both  $\delta_1$  and  $\delta_2$  must equal 1. But this contradicts  $\delta_1 + \delta_2 = 1$ . Hence  $MBL(\mathcal{K})$  has no feasible solution.  $\square$

### 3.3 Necessary Condition for Checking Consistency

In the following, given a PST KB  $\mathcal{K}$ , we introduce a set  $NC(\mathcal{K})$  of linear inequalities such that if  $\mathcal{K}$  is consistent then there is a feasible solution of  $NC(\mathcal{K})$ . That is, the existence of a feasible solution of  $NC(\mathcal{K})$  turns out to be a necessary condition for the consistency of  $\mathcal{K}$ .

As  $MBL(\mathcal{K})$  (see Definition 8),  $NC(\mathcal{K})$  uses rational variables  $v_{id,t,p}$  representing the probability that object  $id$  is at point  $p$  at time  $t$ . As no other kinds of variables are used in the definition of  $NC(\mathcal{K})$ , in this case we obtain a pure system of linear inequalities.

**Definition 9** ( $NC(\mathcal{K})$ ). *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a PST KB.  $NC(\mathcal{K})$  consists of the following (in)equalities:*

$$(1) \quad \forall loc(id, r, t)[\ell, u] \in \mathcal{A}: \quad \ell \leq \sum_{p \in r} v_{id,t,p} \leq u;$$

$$(2) \quad \forall id \in ID, t \in T: \quad \sum_{p \in Space} v_{id,t,p} = 1;$$

$$(3) \quad \forall id \in ID, t \in T, p \in Space: \quad v_{id,t,p} \geq 0;$$

$$(4) \quad \text{for each } f \in \mathcal{F} \text{ (with } k \text{ conjuncts) and } \theta \in \Theta_f \text{ s.t. } \theta(f) = \neg[\bigwedge_{i=1}^k loc(\theta(X_i), \theta(Y_i), \theta(Z_i))],$$

*the inequalities:*

$$\forall p_1 \in \theta(Y_1), p_2 \in \theta(Y_2), \dots, p_k \in \theta(Y_k), \quad \sum_{i=1}^k v_{\theta(X_i), \theta(Z_i), p_i} \leq k - 1.$$

Herein, (in)equalities (1)-(3) are the same as those of  $MBL(\mathcal{K})$  of Definition 8, and have the same meaning. In addition,  $NC(\mathcal{K})$  contains inequalities (4) which impose that, for each ground std-formula of the form  $\neg \bigwedge_{i=1}^k loc(id_i, r_i, t_i)$  and for each  $k$ -tuple of points  $\langle p_1, p_2, \dots, p_k \rangle$  belonging respectively to the regions  $\langle r_1, r_2, \dots, r_k \rangle$ , the sum of the probabilities  $v_{id_i, p_i, t_i}$  that object  $id_i$  is at point  $p_i$  at time  $t_i$ , with  $i \in [1..k]$ , is not greater than  $k - 1$ . As stated in the following theorem, the set consisting of inequalities (4) along with inequalities (1)-(3) turns out to have no feasible solution only if the corresponding PST KB is inconsistent.

**Theorem 3.** *If there is no feasible solution of  $NC(\mathcal{K})$ , then  $\mathcal{K}$  is not consistent.*

*Proof.* Suppose that  $NC(\mathcal{K})$  has no feasible solution. If this is due to the fact there is no feasible solution of  $CC(\langle \mathcal{A}, \emptyset \rangle)$  (introduced in Fact 1), then PST KB  $\langle \mathcal{A}, \emptyset \rangle$  is not consistent, and thus  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  is not consistent as well (since the set  $\mathbf{M}(\langle \mathcal{A}, \mathcal{F} \rangle)$  of the models for  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  is a subset of the set  $\mathbf{M}(\langle \mathcal{A}, \emptyset \rangle)$  of the models for  $\mathcal{K}' = \langle \mathcal{A}, \emptyset \rangle$ ).

Otherwise, there is a feasible solution of  $CC(\langle \mathcal{A}, \emptyset \rangle)$  and thus  $NC(\mathcal{K})$  has no feasible solution due to the fact that at least one of the inequalities in item (4) of Definition 9 is not satisfied by every

solution  $\sigma$  for  $CC(\langle \mathcal{A}, \emptyset \rangle)$ . That is, for each solution  $\sigma$  for  $CC(\langle \mathcal{A}, \emptyset \rangle)$ , there is a ground std-formula  $\theta(f) = \neg \bigwedge_{i=1}^k loc(\theta(X_i), \theta(Y_i), \theta(Z_i))$  such that there exist  $p_1 \in \theta(Y_1), p_2 \in \theta(Y_2), \dots, p_k \in \theta(Y_k)$ , for which  $\sum_{i=1}^k \sigma(v_{\theta(X_i), \theta(Z_i), p_i}) > k - 1$ .

Since every model  $M'$  for  $\mathcal{K}' = \langle \mathcal{A}, \emptyset \rangle$  corresponds to a solution  $\sigma$  for  $CC(\langle \mathcal{A}, \emptyset \rangle)$  such that the sum of the probabilities assigned by  $M'$  to the worlds satisfying an st-atom  $loc(id, \{p\}, t)$  is equal to  $\sigma(v_{id, t, p})$ , the fact that  $\sum_{i=1}^k \sigma(v_{\theta(X_i), \theta(Z_i), p_i}) > k - 1$  holds entails that, for each model  $M'$  for  $\mathcal{K}'$ , the sum of the probabilities assigned by  $M'$  to the worlds satisfying at least one of the st-atoms in  $\theta(f)$  is greater than  $k - 1$ , that is,

$$\forall M' \in \mathbf{M}(\mathcal{K}'), \sum_{i=1}^k \sum_{w \models loc(\theta(X_i), \theta(Y_i), \theta(Z_i))} M'(w) > k - 1 \quad (1)$$

We now recall and use the following well-known result on the minimum probability of a conjunction of probabilistic events among which no correlation is known. Given  $n$  probabilistic events  $e_1, \dots, e_n$  whose (marginal) probabilities are  $\Pr(e_1), \dots, \Pr(e_n)$  respectively,  $\Pr(e_1 \wedge \dots \wedge e_n) \geq \max(0, 1 - n + \sum_{i=1}^n \Pr(e_i))$ . This is one of the Fréchet inequalities (the other one provides an upper bound on the maximum probability) and is implicitly reported already in Boole's work (1854).

In our setting, viewing st-atoms as probabilistic events, the Fréchet inequality entails that for each model  $M'$  for  $\mathcal{K}'$ , the probability that a set  $S$  of st-atoms together satisfy a world is greater than or equal to the maximum between zero and  $1 - |S|$  plus the sum of the probabilities of each st-atom  $a \in S$  according to  $M'$ . That is,

$$\sum_{\substack{w \in \mathcal{W}(\mathcal{K}') \\ \forall a \in S, w \models a}} M'(w) \geq \max \left( 0, 1 - |S| + \sum_{a \in S} \sum_{\substack{w \in \mathcal{W}(\mathcal{K}') \\ w \models a}} M'(w) \right) \quad (2)$$

Equation (1) entails that for each model  $M'$  for  $\mathcal{K}'$ , the term on the right-hand side of Equation (2) evaluates to a value greater than zero when the set  $S$  of st-atoms of  $\theta(f)$  is considered. Therefore, for each model  $M'$  for  $\mathcal{K}'$ , the sum of the probabilities of the worlds satisfying all the st-atoms in  $\theta(f)$  is greater than zero. Since the set  $\mathbf{M}(\langle \mathcal{A}, \mathcal{F} \rangle)$  of the models for  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  is a subset of the set  $\mathbf{M}(\langle \mathcal{A}, \emptyset \rangle)$  of the models for  $\mathcal{K}' = \langle \mathcal{A}, \emptyset \rangle$ , this property also holds for any model for  $\mathcal{K}'$ . Thus, no such  $M'$  can be a model of  $\mathcal{K}$ . Hence  $\mathcal{K}$  is inconsistent.  $\square$

An example of the usage of Theorem 3 is given below, where by checking that  $NC(\mathcal{K})$  has no feasible solution we conclude that the PST KB  $\mathcal{K}$  is not consistent.

**Example 13.** Let  $ID = \{id\}$ ,  $T = [0, 1, 2]$ , and  $Space = \{p_0, p_1, p_2\}$ . Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  such that

- $\mathcal{A}$  consists of the PST atoms  $loc(id, \{p_0\}, 0)[0.4, 1]$ ,  $loc(id, \{p_1\}, 0)[0.5, 1]$ ,  $loc(id, \{p_0\}, 1)[0.8, 1]$ ,  $loc(id, \{p_0\}, 2)[0.8, 1]$ , meaning that,  $id$  is at  $p_0$  and at  $p_1$  at time 0 with probability greater than or equal to 0.4 and 0.5, respectively, and it is at  $p_0$  at times 1 and 2 with probability greater than or equal to 0.8.
- $\mathcal{F}$  consists of the std-formula:  $\neg[loc(id, \{p_0, p_1\}, 0) \wedge loc(id, \{p_0\}, 1) \wedge loc(id, \{p_0\}, 2)]$ , saying that  $id$  cannot be at  $p_0$  at both times 1 and 2 if it was in the region consisting of the points  $p_0$  and  $p_1$  at time point 0.

It is easy to see that  $NC(\mathcal{K})$  contains, among others, the following inequalities:

$$0.4 \leq v_{id,0,p_0} \leq 1$$

$$0.5 \leq v_{id,0,p_1} \leq 1$$

$$0.8 \leq v_{id,1,p_0} \leq 1$$

$$0.8 \leq v_{id,2,p_0} \leq 1$$

$$v_{id,0,p_0} + v_{id,1,p_0} + v_{id,2,p_0} \leq 2$$

$$v_{id,0,p_1} + v_{id,1,p_0} + v_{id,2,p_0} \leq 2$$

where the last two inequalities derive from item (4) of Definition 9. Clearly, these inequalities cannot be satisfied given that  $0.5 \leq v_{id,0,p_1}$ ,  $0.8 \leq v_{id,1,p_0}$ , and  $0.8 \leq v_{id,2,p_0}$ . Thus, since  $NC(\mathcal{K})$  has no feasible solution, we can conclude that  $\mathcal{K}$  is inconsistent.  $\square$

However, we cannot say anything about the consistency of  $\mathcal{K}$  if there is a feasible solution of  $NC(\mathcal{K})$ . The following examples show a case where  $NC(\mathcal{K})$  has a feasible solution and  $\mathcal{K}$  is consistent, and a case where  $NC(\mathcal{K})$  has a feasible solution but  $\mathcal{K}$  is inconsistent.

**Example 14.** Consider the PST KB of Example 13 where  $\mathcal{A}$  is modified such that the probability that  $id$  is at  $p_0$  at times 1 and 2 is greater than or equal to 0.5, instead of 0.8. It is easy to see that there is a feasible solution of  $NC(\mathcal{K})$  in this case. Theorem 3 cannot be used to decide whether  $\mathcal{K}$  is consistent or not. As a matter of fact,  $\mathcal{K}$  is consistent as shown in what follows. Let  $w_1, w_2$  and  $w_3$  be worlds for  $\mathcal{K}$  such that

$$w_1(id, 0) = p_1, w_1(id, 1) = p_1, w_1(id, 2) = p_0,$$

$$w_2(id, 0) = p_0, w_2(id, 1) = p_0, w_2(id, 2) = p_1,$$

$$w_3(id, 0) = p_2, w_3(id, 1) = p_0, w_3(id, 2) = p_2,$$

and let  $M$  be the PDF over  $\mathcal{W}(\mathcal{K})$  such that  $M(w_1) = 0.5$ ,  $M(w_2) = 0.4$ ,  $M(w_3) = 0.1$ , and  $M(w) = 0$  for any other world  $w \in \mathcal{W}(\mathcal{K})$ . It is straightforward to check that  $M$  is a model for  $\mathcal{K}$ .  $\square$

**Example 15.** Again, let  $ID = \{id\}$ ,  $T = [0, 1, 2]$ , and  $Space = \{p_0, p_1, p_2\}$ . Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be such that  $\mathcal{A} = \{loc(id, \{p_0\}, 0)[0.5, 1], loc(id, \{p_0\}, 1)[0.5, 1], loc(id, \{p_0\}, 2)[0.5, 1]\}$ , i.e.,  $id$  is at  $p_0$  at any time in  $T$  with a probability greater than 0.5, and  $\mathcal{F}$  consists of the std-formulas:

$$\neg[loc(id, \{p_0\}, 0) \wedge loc(id, \{p_0\}, 1)],$$

$$\neg[loc(id, \{p_0\}, 1) \wedge loc(id, \{p_0\}, 2)], \text{ and}$$

$$\neg[loc(id, \{p_0\}, 0) \wedge loc(id, \{p_0\}, 2)],$$

saying that  $id$  cannot be at  $p_0$  at time 1 or 2 if it was already there at any previous time value.

A solution for  $NC(\mathcal{K})$  is obtained by assigning the variables  $v_{id,t,p_0}$  and  $v_{id,t,p_1}$  the value 0.5 (where  $t \in [0, 2]$ ), and the other variables the value 0. As  $NC(\mathcal{K})$  has a feasible solution, Theorem 3 says nothing about the fact that  $\mathcal{K}$  is consistent or not. However, it can be checked that  $\mathcal{K}$  is not consistent. Let  $P_t$  be the set of worlds in  $\mathcal{W}(\mathcal{K})$  placing  $id$  at  $p_0$  at time  $t$ , with  $t \in T$ . The std-formulas in  $\mathcal{F}$  entail that every world belonging to two of these three sets must be assigned probability equal to 0 by any model for  $\mathcal{K}$ . That is,  $\mathcal{F}$  entails that  $P_0, P_1, P_2$  are pairwise disjoint sets if we consider only the worlds that can be assigned a probability greater than zero by any model. Now observe that the PST atoms in  $\mathcal{A}$  require that the sum of the probabilities of the worlds in each of these three sets is at least equal to 0.5. Therefore, the overall sum of the probabilities assigned to the worlds in all of these three sets would be greater than or equal to 1.5, which entails that there is no model for  $\mathcal{K}$ .  $\square$

Theorem 3 shows that the consistency of  $\mathcal{K}$  implies the the existence of a feasible solution of  $NC(\mathcal{K})$  and the previous examples show that this is the only relationship between these two

concepts. However, we can make a stronger statement in a special case. For PST KBs where  $\mathcal{F}$  consists of a single ground std-formula (of any arity) constructed from st-atoms referring to single-point regions,  $NC(\mathcal{K})$  has a feasible solution iff the KB is consistent.

**Theorem 4.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  where  $\mathcal{F} = \{\neg[loc(id_1, \{p_1\}, t_1) \wedge \dots \wedge loc(id_n, \{p_n\}, t_n)]\}$ . Then,  $\mathcal{K}$  is consistent iff there is a feasible solution of  $NC(\mathcal{K})$ .*

*Proof.* Using Theorem 3 we need to prove that if there is a feasible solution of  $NC(\mathcal{K})$  then  $\mathcal{K}$  is consistent.

Let  $f$  be the std-formula in  $\mathcal{F}$ . We first observe that if  $f$  contains a pair of st-atoms referring to the same id and time value but different points in  $Space$ , then  $\mathcal{K}$  is consistent. Indeed, in this case,  $f$  is satisfied by every world, and thus it does not impose any restriction on the interpretations for  $\mathcal{K}$  (see Definition 4). Given this, in the following w.l.o.g. we assume that for every distinct pair of st-atoms  $loc(id_i, \{p_i\}, t_i), loc(id_j, \{p_j\}, t_j)$  in  $f$ , if  $id_i = id_j$  then  $t_i \neq t_j$ .

Let  $\sigma$  be a solution for  $NC(\mathcal{K})$ , and  $\mathcal{K}' = \langle \mathcal{A}, \emptyset \rangle$ . Then,  $\sigma$  is a solution for  $CC(\mathcal{K}')$  (from Fact 1) and corresponds to a model  $M'$  for  $\mathcal{K}'$  such that (i) the sum of the probabilities assigned by  $M'$  to the worlds satisfying an st-atom  $loc(id, \{p\}, t)$  (i.e., the marginal probability of  $loc(id, \{p\}, t)$ ) is equal to  $\sigma(v_{id,t,p})$ , and (ii)  $\sum_{i=1}^k \sigma(v_{id_i,t_i,p_i}) \leq k - 1$ . Viewing each st-atom  $loc(id_i, \{p_i\}, t_i)$  in  $f$  as a probabilistic event whose (marginal) probability is  $\sigma(v_{id_i,t_i,p_i})$ , the Fréchet inequality (recalled in the proof of Theorem 3) entails that the minimum probability that the st-atoms in  $f$  occur all together is equal to  $\max(0, 1 - k + \sum_{i=1}^k \sigma(v_{id_i,t_i,p_i}))$ . This is equal to zero since the second argument of function  $\max$  is not greater than zero due to the fact that (ii) holds. The fact that the minimum probability that the st-atoms in  $f$  simultaneously occur is equal to zero suffices to ensure that there is at least one model  $M''$  for  $\mathcal{K}'$  such that  $M''$  assigns probability equal to zero to all the worlds for  $\mathcal{K}'$  that do not satisfy  $f$ . As  $f$  is the only std-formula in  $\mathcal{F}$ , and every world not satisfying  $f$  is assigned a probability equal to zero by  $M''$ , it follows that  $M''$  is a model for  $\mathcal{K}$  too.  $\square$

The following example shows that considering even a binary std-formula containing an st-atom referring to a region consisting of two points, it may happen that there is a feasible solution of  $NC(\mathcal{K})$  even if  $\mathcal{K}$  is not consistent.

**Example 16.** Let  $ID = \{id\}$ ,  $T = [0, 1]$ , and  $Space = \{p_0, p_1, p_2\}$ . Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be such that  $\mathcal{A} = \{loc(id, \{p_0\}, 0)[0.4, 1], loc(id, \{p_1\}, 0)[0.4, 1], loc(id, \{p_0\}, 1)[0.4, 1], \}$ , and  $\mathcal{F}$  is the std-formula  $\neg[loc(id, \{p_0, p_1\}, 0) \wedge loc(id, \{p_2\}, 1)]$  saying that  $id$  cannot move to point  $p_2$  at time 1 if it was in either  $p_0$  or  $p_1$  at time 0. It is easy to see that  $NC(\mathcal{K})$  is feasible but  $\mathcal{K}$  is not consistent.  $\square$

In Section 3.5, we will present a method for deciding in polynomial time the consistency of PST KBs where binary std-formulas satisfying some ‘‘acyclicity’’ conditions are used. It turns out that the consistency of the PST KBs of both Examples 15 and 16 can be decided in polynomial time using our approach.

### 3.4 Unary Std-Formulas are Tractable

We start by identifying a tractable case of the consistency checking problem: when all std-formulas are *unary*, that is, each formula in  $\mathcal{F}$  consists of only one st-atom and possibly a conjunction of built-in predicates (i.e., in Definition 2,  $k = 1$ ).

**Example 17.** The constraint “*there is no object in region  $r$  at any time between 5 and 10*” can be expressed by the following unary std-formula:  $\forall X_1, Z_1 \neg[loc(X_1, r, Z_1) \wedge Z_1 \geq 5 \wedge 10 \geq Z_1]$ .

The constraint “*object  $id$  is always in region  $r$* ” can be expressed as:

$$\forall Y_1, Z_1 \neg[loc(id, Y_1, Z_1) \wedge Y_1 \text{nov } r].^7 \quad \square$$

The following theorem states that checking consistency is tractable if only *unary* std-formulas are considered.

**Theorem 5.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a PST KB such that  $\mathcal{F}$  consists of unary std-formulas only. Then, deciding whether  $\mathcal{K}$  is consistent is in PTIME.*

*Proof.* The statement follows from the fact that if  $\mathcal{F}$  consists of *unary* std-formulas only,  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  is equivalent to (i.e., it has exactly the same set of models as)  $\mathcal{K}' = \langle \mathcal{A}', \emptyset \rangle$ , where  $\mathcal{A}'$  consists of the atoms in  $\mathcal{A}$  plus the atom  $loc(\theta(X_i), \theta(Y_i), \theta(Z_i))[0, 0]$  for each ground std-formula  $\theta(f) = \neg[loc(\theta(X_i), \theta(Y_i), \theta(Z_i))]$ , where  $f \in \mathcal{F}$  and  $\theta \in \Theta_f$ . Since,  $\bigcup_{f \in \mathcal{F}} \Theta_f$  is polynomial w.r.t. the size of  $\mathcal{K}$ , the size of  $\mathcal{A}'$  (and thus of  $\mathcal{K}'$ ) increases by a polynomial number of atoms. Hence, we can apply Fact 1, which entails that the consistency of PST KBs with  $\mathcal{F} = \emptyset$  can be decided in PTIME.  $\square$

### 3.5 Tractable Binary Std-Formulas

In the following we focus on PST KBs where all std-formulas are binary. This is a restricted but expressive class of std-formulas that allow us to impose several practical constraints in many application contexts. As a matter of fact, both  $f_1$  and  $f_2$  of Example 2 as well as  $f_4$  of Example 4 are binary std-formulas. Furthermore, using the approach suggested in the proof of Theorem 5, we can assume that unary std-formulas are encoded as PST atoms. Thus, all the results stated in this section straightforwardly apply to the case where both unary and binary std-formulas are in the PST KB.

We start by noting that consistency checking was proved to be feasible by Grant et al. (2010) for the case where reachability definitions (but no other integrity constraints) were allowed. As we showed in Example 3 a reachability definition can be expressed by a binary std-formula. Hence, in the special case where all the binary std-formulas represent reachability definitions, consistency checking is tractable.

In the general case of a PST KB  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  such that  $\mathcal{F}$  consists of binary std-formulas, we define an undirected graph, called the *std-graph*, where each maximal independent set represents a world for  $\mathcal{K}$  satisfying all the std-formulas in  $\mathcal{F}$ .<sup>8</sup> We will later use this graph to characterize binary std-formulas for which the consistency checking problem turns out to be tractable.

**Definition 10** (std-graph). *Given a PST KB  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  such that  $\mathcal{F}$  consists of binary std-formulas, the std-graph  $\mathcal{G}(\mathcal{K})$  is an undirected graph  $\langle V, E \rangle$  whose sets of vertices  $V$  and edges  $E$  are such that:*

- 1)  $V$  consists of the set of all ground st-atoms of the form  $loc(id, \{p\}, t)$  where  $id \in ID$ ,  $p \in Space$ , and  $t \in T$ ;

7. Recall that every substitution for  $Y_1$  must be a region containing a single point.

8. A maximal independent set in the std-graph is a maximal independent set in an undirected graph whose vertices are st-atoms. The formal relationship between maximal independent sets in the std-graph and worlds of a PST KB is given in Proposition 1.

2)  $E$  consists of

- i) an edge between every pair of ground st-atoms in  $V$  referring to the same object and time value, that is, for each  $p_1, p_2 \in \text{Space}$ ,  $p_1 \neq p_2$ ,  $id \in ID$ ,  $t \in T$ ,  $\langle \text{loc}(id, \{p_1\}, t), \text{loc}(id, \{p_2\}, t) \rangle \in E$ ;
- ii) an edge  $\langle \text{loc}(id_1, \{p_1\}, t_1), \text{loc}(id_2, \{p_2\}, t_2) \rangle$  between every pair of ground st-atoms in  $V$  such that  $\neg[\text{loc}(id_1, r_1, t_1) \wedge \text{loc}(id_2, r_2, t_2)]$  with  $p_1 \in r_1$  and  $p_2 \in r_2$  belongs to the set of ground std-formulas that can be derived from  $\mathcal{F}$ .

We will just write  $\mathcal{G}$  instead of  $\mathcal{G}(\mathcal{K})$  where  $\mathcal{K}$  is known.

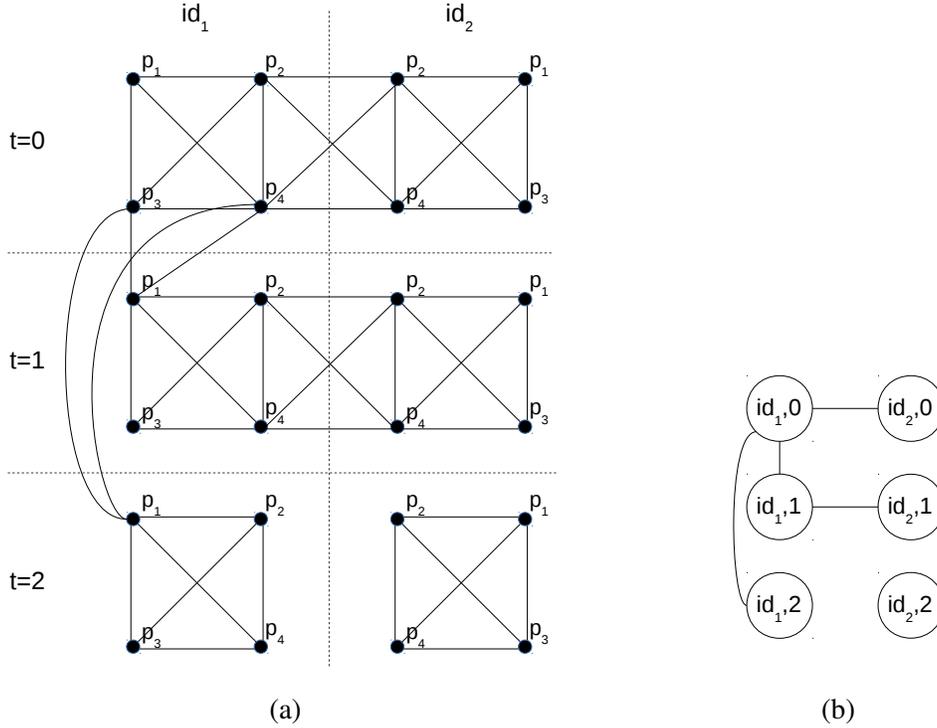
Basically, each edge of  $\mathcal{G}$  connects a pair of st-atoms that cannot belong together in any world satisfying the std-formulas in  $\mathcal{F}$ . In particular, an edge of type *i*) connects two st-atoms representing the fact that an object is in two places at the same time — this is not admitted according to the definition of world (see Definition 4). An edge of type *ii*) connects two st-atoms representing a fact (i.e., object  $id_1$  is at point  $p_1$  at time  $t_1$  and object  $id_2$  is at point  $p_2$  at time  $t_2$ ) not consistent with a ground std-formula entailed by  $\mathcal{F}$ .

The structure of  $\mathcal{G}$  is as follows. For each  $\langle id, t \rangle$  pair,  $\mathcal{G}$  contains a clique<sup>9</sup> of size  $|\text{Space}|$  which consists of a vertex  $\text{loc}(id, \{p\}, t)$  for each point in  $\text{Space}$  and edges of type *i*) — in the following, we refer to this clique as the *clique for  $\langle id, t \rangle$  pair* or the  *$\langle id, t \rangle$  clique*. For each ground std-formula  $f = \neg[\text{loc}(id_1, r_1, t_1) \wedge \text{loc}(id_2, r_2, t_2)]$  that can be derived from  $\mathcal{F}$ ,  $\mathcal{G}$  contains a clique of size  $|r_1| + |r_2|$  which consists of a vertex  $\text{loc}(id_1, \{p_1\}, t_1)$  for each point in  $r_1$ , a vertex  $\text{loc}(id_2, \{p_2\}, t_2)$  for each point in  $r_2$ , and edges of both types *i*) and *ii*) — we refer to this clique as the *clique for std-formula  $f$* .

**Example 18.** Let  $ID = \{id_1, id_2\}$ ,  $T = [0, 1, 2]$ , and  $\text{Space} = \{p_1, p_2, p_3, p_4\}$ . Assume that  $\mathcal{F}$  consists of the following std-formulas:

- $f_1 = \forall X_1, X_2, Z_1 \neg[\text{loc}(X_1, \{p_2, p_4\}, Z_1) \wedge \text{loc}(X_2, \{p_2, p_4\}, Z_1) \wedge X_1 \neq X_2 \wedge 0 \leq Z_1 \wedge Z_1 \leq 1]$ , saying that there cannot be two distinct objects in the region consisting of points  $\{p_2, p_4\}$  at times 0 and 1;
- $f_2 = \forall Z_1 \neg[\text{loc}(id_1, \{p_3, p_4\}, 0) \wedge \text{loc}(id_1, \{p_1\}, Z_1) \wedge 1 \leq Z_1 \wedge Z_1 \leq 2]$ , saying that object  $id_1$  cannot reach point  $p_1$  starting from region  $\{p_3, p_4\}$  at time 0 in 1 or 2 time units.

The std-graph  $\mathcal{G}$  is shown in Figure 2(a), where, for the sake of readability, vertices are labelled with the names of the points in  $\text{Space}$  to which they refer, while the id and time value of each vertex is reported on the column and the row to which it belongs (for instance, vertex  $\text{loc}(id_1, \{p_1\}, 0)$  is represented by the circle in the upper-left corner). Observe that  $\mathcal{G}$  consists of 10 (maximal) cliques, one for each of the 6  $\langle id, t \rangle$  pairs, and one for each of the 4 ground std-formulas derived from  $\mathcal{F}$ . Specifically, each  $\langle id, t \rangle$ -clique, with  $id \in \{id_1, id_2\}$  and  $t \in [1..2]$ , consists of the four vertices  $\text{loc}(id, \{p_k\}, t)$  with  $k \in [1..4]$ , while the cliques for the ground std-formulas are the following sets:  $\{\text{loc}(id_1, \{p_2\}, 0), \text{loc}(id_2, \{p_2\}, 0), \text{loc}(id_1, \{p_4\}, 0), \text{loc}(id_2, \{p_4\}, 0)\}$ ,  $\{\text{loc}(id_1, \{p_2\}, 1), \text{loc}(id_2, \{p_2\}, 1), \text{loc}(id_1, \{p_4\}, 1), \text{loc}(id_2, \{p_4\}, 1)\}$ ,  $\{\text{loc}(id_1, \{p_3\}, 0), \text{loc}(id_1, \{p_4\}, 0), \text{loc}(id_1, \{p_1\}, 1)\}$ , and  $\{\text{loc}(id_1, \{p_3\}, 0), \text{loc}(id_1, \{p_4\}, 0), \text{loc}(id_1, \{p_1\}, 2)\}$ . □


 Figure 2: (a) Std-graph  $\mathcal{G}$ ; (b) Auxiliary-graph  $\mathcal{AG}$ .

It is worth noting that the cliques for the  $\langle id, t \rangle$  pairs as well as those for the std-formulas are not maximal cliques in general, as shown in the following.

**Example 19.** Continuing Example 18, assume that  $\mathcal{F}$  is augmented with the following (ground) std-formulas:

- $f_3 = \neg[loc(id_2, \{p_3\}, 1) \wedge loc(id_2, \{p_2\}, 2)]$
- $f_4 = \neg[loc(id_2, \{p_4\}, 1) \wedge loc(id_2, \{p_2\}, 2)]$

Thus, the clique for std-formula  $f_3$  consists of the set  $\{loc(id_2, \{p_3\}, 1), loc(id_2, \{p_2\}, 2)\}$ , while that for  $f_4$  consists of the set  $\{loc(id_2, \{p_4\}, 1), loc(id_2, \{p_2\}, 2)\}$ . Both these cliques are not maximal ones, as they are included in the clique consisting of the set of vertices  $\{loc(id_2, \{p_3\}, 1), loc(id_2, \{p_4\}, 1), loc(id_2, \{p_2\}, 2)\}$ . This is basically due to the fact that the constraint imposed by  $f_3$  and  $f_4$  can be expressed more succinctly, as  $\neg[loc(id_2, \{p_3, p_4\}, 1) \wedge loc(id_2, \{p_2\}, 2)]$  whose associated clique is maximal.  $\square$

The following proposition follows from the definition of std-graph and the fact that every object must be at a unique point for each time value.

**Proposition 1.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a PST KB where  $\mathcal{F}$  consists of binary std-formulas only. Every maximal independent set of  $\mathcal{G}$  consisting of a vertex from each  $\langle id, t \rangle$  clique is in a one-to-one correspondence with a world  $w \in \mathcal{W}$  such that  $w \models \mathcal{F}$ .*

9. Note that we use the terminology *clique* for a complete subgraph of  $\mathcal{G}$ , and *maximal clique* for a clique not contained in any other clique. We point this out since maximal cliques are often called simply “cliques”.

Observe that if there is no maximal independent set of  $\mathcal{G}$  having the property stated in Proposition 1, then  $\mathcal{G}$  must contain a maximal clique including at least two  $\langle id, t \rangle$  cliques (this happens for instance if  $\mathcal{F}$  contains an std-formula  $\neg[loc(id, Space, t_1) \wedge loc(id, Space, t_2)]$ ). In this case, no world satisfies  $\mathcal{F}$  and the PST KB is trivially inconsistent. In the general case, there may be an exponential number of maximal independent sets of  $\mathcal{G}$  having the property stated in Proposition 1, and the fact that the PST KB is consistent or not can be decided using  $\mathcal{G}$  as explained in what follows.

### 3.5.1 CLIQUE-ACYCLIC STD-GRAPHS

Our characterization of tractable cases of the consistency checking problem focuses on KBs with binary std-formulas such that their std-graphs have a property, that we call *clique-acyclic*. We start with some preliminary definitions.

**Definition 11** (Binary maximal clique). *For std-graph  $\mathcal{G}$  we call a clique binary maximal iff it contains vertices from two  $\langle id, t \rangle$  pairs and is not properly included in any clique that contains vertices from the same  $\langle id, t \rangle$  pairs.*

In particular, the std-graph  $\mathcal{G}$  of Figure 2(a) has 4 binary maximal cliques all involving  $id_1$ .

**Definition 12** (Clique-subgraph). *We call a subgraph of std-graph  $\mathcal{G}$  a clique subgraph iff it contains all the vertices of  $\mathcal{G}$ , one edge from each  $\langle id, t \rangle$  clique as well as one new distinct edge from each binary maximal clique.*

So any clique-subgraph of the std-graph of the graph  $\mathcal{G}$  of Figure 2(a) has 10 edges, one from each of the 6  $\langle id, t \rangle$  pairs and one from each binary maximal clique.

**Definition 13** (Clique-acyclic std-graph). *Std-graph  $\mathcal{G}$  is said to be clique-acyclic iff all of its clique-subgraphs are acyclic graphs (that is, forests).  $\mathcal{G}$  is called clique-cyclic if it is not clique-acyclic.*

Basically, clique-acyclicity means that no cycle can be found in the std-graph after compressing all binary maximal cliques into single edges and using just one edge of each  $\langle id, t \rangle$ -clique. It is easy to see that the std-graph shown in Figure 2(a) is clique-acyclic. As stated in the next proposition, clique-acyclicity can be checked using the following auxiliary graph which basically compresses a clique-subgraph to its essential structure. It will be clear from the definition that the same auxiliary graph is obtained from all clique-subgraphs of a graph  $\mathcal{G}$ .

**Definition 14** (Auxiliary graph). *The auxiliary graph for  $\mathcal{G}$  is the undirected graph  $\mathcal{AG} = \langle V', E' \rangle$  such that:*

- $V'$  consists of a vertex for each  $\langle id, t \rangle$  pair, with  $id \in ID$  and  $t \in T$ ;
- $E'$  consists of all the edges for binary maximal cliques in a clique-subgraph where each previous vertex  $loc(id, \{p\}, t)$  is replaced by the corresponding  $\langle id, t \rangle$  pair. We will denote as  $C(e)$  the binary maximal clique  $C$  from which the edge  $e$  originated.

The auxiliary graph for the std-graph of Figure 2(a) is shown in Figure 2(b). As an example,  $C(e)$  for edge  $e = \langle (id_1, 0), (id_1, 1) \rangle$  consists of the following set of vertices  $\{loc(id_1, \{p_3\}, 0), loc(id_1, \{p_4\}, 0), loc(id_1, \{p_1\}, 1)\}$ .

The following proposition follows from Definitions 13 and 14.

**Proposition 2.** *Std-graph  $\mathcal{G}$  is clique-acyclic iff the auxiliary graph  $\mathcal{AG}$  is acyclic (that is,  $\mathcal{AG}$  is a forest).*

In the following, we introduce a set of linear inequalities that can be used to check the consistency of PST KBs where the std-formulas are binary and the auxiliary graph is acyclic. Then in the next subsection we will do the same for a special case of a cyclic auxiliary graph. In both cases we will be working with a single connected component of the auxiliary graph. Suppose that  $\mathcal{AG}$  has  $n$  connected components  $C_1, \dots, C_n$  and let  $\mathcal{K}_i$  be the PST KB corresponding to  $C_i$  for each  $i$ . If  $\mathcal{K}$  has a model  $M$  we show how to obtain a model  $M_i$  for  $\mathcal{K}_i$ . Let  $w_i$  be a world appropriate to the  $\langle id, t \rangle$  pairs for  $\mathcal{K}_i$ . Let  $W_i$  be all the worlds of  $\mathcal{K}$  that extend  $w_i$  to the  $\langle id, t \rangle$  values not in  $\mathcal{K}_i$ . Define  $M_i(w_i) = \sum_{w \in W_i} M(w)$ . Going in the other direction, suppose that  $M_1, \dots, M_n$  are models for the PST KBs corresponding to  $C_i, \dots, C_n$  respectively. For any world  $w$  for  $\mathcal{K}$ , let  $w_1, \dots, w_n$  be the restrictions of  $w$  to  $C_1, \dots, C_n$  respectively. Defining  $M(w) = \prod_{i=1}^n M_i(w_i)$  is then a model of  $\mathcal{K}$ . We have shown the following result.

**Proposition 3.**  *$\mathcal{K}$  is consistent iff all of the PST KBs corresponding to the connected components of  $\mathcal{G}$  (and hence of  $\mathcal{AG}$ ) are consistent.*

Hence in our proofs it suffices to assume that  $\mathcal{AG}$  has a single connected component.<sup>10</sup>

**Definition 15** ( $BC(\mathcal{K})$ ). *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a PST KB such that  $\mathcal{F}$  consists of binary std-formulas only.  $BC(\mathcal{K})$  consists of the following (in)equalities:*

$$(1) \quad \forall loc(id, r, t)[\ell, u] \in \mathcal{A}: \quad \ell \leq \sum_{p \in r} v_{id,t,p} \leq u;$$

$$(2) \quad \forall id \in ID, t \in T: \quad \sum_{p \in Space} v_{id,t,p} = 1;$$

$$(3) \quad \forall id \in ID, t \in T, p \in Space: \quad v_{id,t,p} \geq 0;$$

$$(4) \quad \text{For each edge } e \text{ of } \mathcal{AG}: \quad \sum_{loc(id, \{p\}, t) \in C(e)} v_{id,t,p} \leq 1.$$

The following theorem states that checking whether  $BC(\mathcal{K})$  has a feasible solution is equivalent to deciding the consistency of PST KBs  $\mathcal{K}$  where std-formulas are binary and generate an acyclic auxiliary graph. The intuition behind this result is as follows. If there is a feasible solution of  $BC(\mathcal{K})$ , then inequalities (1)-(3) entail that for each  $\langle id, t \rangle$  pair, the events that object  $id$  at time  $t$  is at point  $p \in Space$  can be arranged to fit the whole probability space without overlapping. Roughly speaking this would suffice to define a model for a PST KB without std-formulas by combining the distributions obtained for each  $\langle id, t \rangle$  pair. Considering std-formulas means that some events cannot occur together, that is, they cannot coexist in the same portion of the probability space. Intuitively, the fact that inequality (4) of  $BC(\mathcal{K})$  is satisfied for an edge  $e$  entails that the events corresponding to the st-atoms in  $C(e)$  can be arranged in distinct portions of the probability space (avoiding overlaps). If the auxiliary graph  $\mathcal{AG}$  is acyclic, this reasoning can be inductively repeated for each edge of  $\mathcal{AG}$  without ever reconsidering the arrangement of events corresponding to the st-atoms considered previously. The satisfiability of inequality (4) of  $BC(\mathcal{K})$  is also a necessary

10. We do not indicate this in the statement of the following theorem because if  $\mathcal{AG}$  has several connected components, the proof works for each one and then Proposition 3 is applied.

condition for the consistency of  $\mathcal{K}$  as the fact that it is not satisfied for an edge  $e$  of  $\mathcal{AG}$  intuitively means that the events corresponding to the st-atoms in  $C(e)$ , even taken alone, cannot be arranged in the probability space without overlapping.

**Theorem 6.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a PST KB where  $\mathcal{F}$  consists of binary std-formulas only. If  $\mathcal{AG}$  is acyclic, then  $\mathcal{K}$  is consistent iff there is a feasible solution of  $BC(\mathcal{K})$ .*

*Proof.* We will use the notation  $\mathcal{K}'$  for  $\mathcal{K}$  with  $\mathcal{F}$  removed, that is,  $\mathcal{K}' = \langle \mathcal{A}, \emptyset \rangle$ . Then,  $\mathbf{M}(\mathcal{K})$  is a subset of  $\mathbf{M}(\mathcal{K}')$ .

( $\Rightarrow$ ) We prove the contrapositive: if there is no feasible solution of  $BC(\mathcal{K})$ , then  $\mathcal{K}$  is inconsistent. The simplest case is where  $BC(\mathcal{K})$  has no feasible solution due to the fact that  $BC(\mathcal{K}')$  has no feasible solution. Then there are no inequalities (4), hence  $\mathcal{K}'$  is not consistent, and thus  $\mathcal{K}$  is not consistent as well.

Otherwise, there is a feasible solution of  $BC(\mathcal{K}')$  and thus  $BC(\mathcal{K})$  has no feasible solution due to the fact that at least one of the inequalities of item (4) of Definition 15 is not satisfied. Therefore, there is an edge  $e$  of  $\mathcal{AG}$  such that  $\sum_{loc(id, \{p\}, t) \in C(e)} v_{id, t, p} > 1$ . To show that  $\mathcal{K}$  is inconsistent we show that no interpretation can be a model of  $\mathcal{K}$ . So suppose that  $M$  is an interpretation that is a model for  $\mathcal{K}$ . Since a pair  $A_1, A_2$  of st-atoms belong to  $C(e)$  iff (a)  $\neg[A_1 \wedge A_2]$  belongs to the set of ground formulas that can be derived by  $\mathcal{F}$ , or (b)  $A_1 = loc(id, \{p_1\}, t)$  and  $A_2 = loc(id, \{p_2\}, t)$  where  $p_2 \neq p_1$ , there is no world  $w \in \mathcal{W}(\mathcal{K})$  such that  $M(w) > 0$ ,  $w \models A_1$ , and  $w \models A_2$  (in particular, in case (b) there is no world  $w \in \mathcal{W}(\mathcal{K})$  such that  $w \models A_1$ , and  $w \models A_2$ , whether or not  $M(w) > 0$ ). That is, every world  $w$  for  $\mathcal{K}$  which is assigned a non-zero probability by  $M$  is such that it satisfies at most one st-atom in  $C(e)$ . Hence,  $\mathcal{W}(\mathcal{K})$  can be partitioned into  $|C(e)| + 2$  pairwise disjoint sets as follows:

- (i) the set of worlds  $w \in \mathcal{W}(\mathcal{K})$  such that  $M(w) = 0$  (this set includes all the worlds satisfying at least two st-atoms in  $C(e)$ )
- (ii) the set of worlds  $w \in \mathcal{W}(\mathcal{K})$  such that  $M(w) > 0$  and  $w$  does not satisfy any st-atom in  $C(e)$ ;
- (iii) for each st-atom  $A_i \in C(e)$ , the set of worlds  $w \in \mathcal{W}(\mathcal{K})$  such that  $M(w) > 0$ ,  $w \models A_i$ , and  $w \not\models A_j$  where  $A_j \in C(e)$  (and is different from  $A_i$ ).

Therefore, the sum of the probabilities of the worlds satisfying at least one st-atom in  $C(e)$  is equal to the sum, for each st-atom  $A$  in  $C(e)$ , of the sum of the probabilities of the worlds satisfying atom  $A \in C(e)$  but no other st-atoms in  $C(e)$ , that is

$$\sum_{\substack{w \in \mathcal{W}(\mathcal{K}) \\ w \models loc(id, \{p\}, t) \\ loc(id, \{p\}, t) \in C(e)}} M(w) = \sum_{loc(id, \{p\}, t) \in C(e)} \sum_{\substack{w \in \mathcal{W}(\mathcal{K}) \\ w \models loc(id, \{p\}, t) \\ w \not\models loc(id', \{p'\}, t') \\ loc(id', \{p'\}, t') \in C(e)}} M(w).$$

Now recall that any model for  $\mathcal{K}$  is also a model for  $\mathcal{K}'$ . The above-mentioned partitioning of  $\mathcal{W}(\mathcal{K})$  also holds for  $\mathcal{W}(\mathcal{K}')$ . Therefore, there is a solution  $\sigma$  for  $BC(\mathcal{K}')$ , one-to-one corresponding to a model  $M$  for  $\mathcal{K}'$ , such that the sum of the probabilities assigned by  $M$  to the worlds for  $\mathcal{K}'$  satisfying an st-atom  $loc(id, \{p\}, t)$  is equal to the value assigned to variable  $v_{id, t, p}$  by  $\sigma$ . In particular,

for a given st-atom  $loc(id, \{p\}, t) \in C(e)$ , we have that

$$\sum_{\substack{w \in \mathcal{W}(\mathcal{K}') \\ w \models loc(id, \{p\}, t)}} M(w) = \sum_{\substack{w \in \mathcal{W}(\mathcal{K}') \\ w \models loc(id, \{p\}, t) \\ loc(id, \{p\}, t) \in C(e) \\ w \not\models loc(id', \{p'\}, t') \\ loc(id', \{p'\}, t') \in C(e)}} M(w) = v_{id,t,p}.$$

Considering all the st-atoms in  $C(e)$  we obtain that

$$\begin{aligned} \sum_{\substack{w \in \mathcal{W}(\mathcal{K}') \\ w \models loc(id, \{p\}, t) \\ loc(id, \{p\}, t) \in C(e)}} M(w) &= \sum_{loc(id, \{p\}, t) \in C(e)} \sum_{\substack{w \in \mathcal{W}(\mathcal{K}') \\ w \models loc(id, \{p\}, t) \\ w \not\models loc(id', \{p'\}, t') \\ loc(id', \{p'\}, t') \in C(e)}} M(w) = \\ &= \sum_{loc(id, \{p\}, t) \in C(e)} v_{id,t,p} > 1. \end{aligned}$$

The inequality above holds since the inequality (4) of Definition 15 is not satisfied for edge  $e$ . Finally, since the following inequality holds due to the definition of model

$$\sum_{w \in \mathcal{W}(\mathcal{K}')} M(w) \geq \sum_{\substack{w \in \mathcal{W}(\mathcal{K}') \\ w \models loc(id, \{p\}, t)}} M(w),$$

and the latter is strictly greater than one, it follows that  $M$  is not a PDF over  $\mathcal{W}(\mathcal{K}')$ , meaning that  $M$  is a not model for  $\mathcal{K}'$ . Hence no interpretation can be a model for  $\mathcal{K}$  and so  $\mathcal{K}$  is inconsistent.

( $\Leftarrow$ ) We now prove that if the auxiliary graph  $\mathcal{AG}$  is acyclic and  $BC(\mathcal{K})$  has a feasible solution, then  $\mathcal{K}$  is consistent. We prove this by mathematical induction on the number of edges of  $\mathcal{AG}$ .

*Base case:* The number of edges is 0, so  $\mathcal{AG}$  consists of a set of isolated vertices. The fact that  $\mathcal{AG}$  contains no edges means that the set of ground std-formulas that can be derived from  $\mathcal{F}$  is empty. Thus there is a feasible solution of  $BC(\mathcal{K})$  iff there is a feasible solution of  $CC(\mathcal{K}')$ . Using Fact 1, we obtain that  $\mathcal{K}$  is consistent.

*Inductive step:* We prove that if the statement holds for  $\alpha$  edges of  $\mathcal{AG}$ , then it holds for  $\alpha + 1$  edges. So we are given an acyclic graph  $\mathcal{AG}$  with  $\alpha + 1$  edges, that we write as  $\mathcal{AG}^{\alpha+1}$ . By acyclicity we can choose a subgraph,  $\mathcal{AG}^\alpha$  with  $\alpha$  edges such that the new edge connects a vertex to an isolated vertex of  $\mathcal{AG}^\alpha$  and  $\mathcal{AG}^\alpha$  is clearly acyclic. Let  $\mathcal{K}^\alpha = \langle \mathcal{A}, \mathcal{F}^\alpha \rangle$  be the PST KB where  $\mathcal{F}^\alpha$  consists of the subset of ground std-formulas derived from  $\mathcal{F}$  so that the std-graph  $\mathcal{G}^\alpha$  has the auxiliary graph  $\mathcal{AG}^\alpha$ . The induction hypothesis is that if there is a feasible solution of  $BC(\mathcal{K}^\alpha)$ , then  $\mathcal{K}^\alpha$  is consistent.

We write  $\mathcal{K}^{\alpha+1} = \langle \mathcal{A}, \mathcal{F}^{\alpha+1} \rangle$  where  $\mathcal{F}^{\alpha+1}$  consists of the ground std-formulas in  $\mathcal{F}^\alpha$  plus those corresponding to the new edge  $e$ , so that  $\mathcal{AG}^{\alpha+1}$  turns out to be the auxiliary graph for  $\mathcal{G}(\mathcal{K}^{\alpha+1})$ . Assuming the induction hypothesis, we now show that if there is a feasible solution of  $BC(\mathcal{K}^{\alpha+1})$ , then  $\mathcal{K}^{\alpha+1}$  is consistent.

Let  $\sigma$  be a solution for  $BC(\mathcal{K}^{\alpha+1})$ . Obviously,  $\sigma$  is also a solution for  $BC(\mathcal{K}^\alpha)$ , as it consists of a subset of the (in)equalities in  $BC(\mathcal{K}^{\alpha+1})$ . The fact that  $\sigma$  is a solution for  $BC(\mathcal{K}^\alpha)$  means that there is a model  $M^\alpha$  for  $\mathcal{K}^\alpha$  such that the sum of the probabilities assigned by  $M^\alpha$  to the worlds for

$\mathcal{K}^\alpha$  satisfying an st-atom  $loc(id, \{p\}, t)$  is equal to the value  $\sigma(v_{id,t,p})$  assigned to variable  $v_{id,t,p}$  by  $\sigma$ . Since  $\sigma$  is a solution also for  $BC(\mathcal{K}^{\alpha+1})$ , and thus inequality (4) of Definition 15 holds, it follows that

$$\sum_{loc(id, \{p\}, t) \in C(e)} \sum_{\substack{w \in \mathcal{W}(\mathcal{K}^\alpha), \\ w \models loc(id, \{p\}, t)}} M^\alpha(w) = \sum_{loc(id, \{p\}, t) \in C(e)} \sigma(v_{id,t,p}) \leq 1.$$

That is,  $M^\alpha$  is a model for  $\mathcal{K}^\alpha$  with the property that the sum of the probabilities of the worlds satisfying at least one st-atom in  $C(e)$  is less than or equal to one.

If  $M^\alpha$  is such that it assigns probability equal to zero to every world in  $\mathcal{W}(\mathcal{K}^\alpha)$  containing at least two st-atoms in  $C(e)$ , then we are done, since  $M^\alpha$  turns out to be a model for  $\mathcal{K}^{\alpha+1}$  too.

Otherwise, starting from  $M^\alpha$ , we show that a model  $M^{\alpha+1}$  for  $\mathcal{K}^{\alpha+1}$  can be constructed by reasoning as follows.

We recall that the new edge  $e$  being added to  $\mathcal{AG}^\alpha$  to obtain  $\mathcal{AG}^{\alpha+1}$  is of the form  $\langle (id_1, t_1), (id_2, t_2) \rangle$ . Let  $C_1(e)$  and  $C_2(e)$  be the sets of st-atoms in  $C(e)$  of the form  $loc(id_1, \{p\}, t_1)$  and  $loc(id_2, \{p\}, t_2)$ , respectively. Hence,  $C(e) = C_1(e) \cup C_2(e)$  and  $C_1(e) \cap C_2(e) = \emptyset$ . Let  $\pi = \langle w_1, w_2, \dots, w_n \rangle$  be a permutation of the worlds  $w \in \mathcal{W}(\mathcal{K}^\alpha)$  which are assigned by  $M^\alpha$  a probability greater than zero (i.e.,  $M^\alpha(w) > 0$ ) and such that the first  $k$  worlds in  $\pi$  satisfy an st-atom in  $C_1(e)$ . Note that,  $\mathcal{W}(\mathcal{K}^\alpha) = \mathcal{W}(\mathcal{K}^{\alpha+1})$ , and that a world in  $\mathcal{W}(\mathcal{K}^\alpha)$  consists of at most one atom in  $C_1(e)$ , as it is a function whose domain is  $ID \times T$  (see Definition 4) and all of the st-atoms in  $C_1(e)$  refer to the same id and time value. Let us denote by  $\pi_i$  the sum of the probabilities of the first  $i$  worlds in  $\pi$ , that is,  $\pi_i = \sum_{j=1}^i M^\alpha(w_j)$ , where  $w_j$  is in  $\pi$ . Thus,  $\pi_k$  is equal to the sum of the probabilities assigned by  $M^\alpha$  to the worlds satisfying an st-atom in  $C_1(e)$ , i.e.,

$$\pi_k = \sum_{\substack{w \in \mathcal{W}(\mathcal{K}^\alpha), \\ w \models loc(id, \{p\}, t), \\ loc(id, \{p\}, t) \in C_1(e)}} M^\alpha(w) = \sum_{loc(id, \{p\}, t) \in C_1(e)} \sigma(v_{id,t,p}).$$

Let  $S$  be the subset of  $Space$  consisting of points  $p$  such that variables  $v_{id_2, t_2, p}$  are assigned a value greater than zero by  $\sigma$ , that is,  $S = \{p \mid p \in Space, \sigma(v_{id_2, t_2, p}) > 0\}$ . Observe that these variables correspond to st-atoms referring to endpoint  $(id_2, t_2)$  of edge  $e$ . Let  $\nu = \langle p_1, p_2, \dots, p_m \rangle$  be a permutation of the points in  $S$  such that the first  $h$  points  $p$  correspond to the variables  $v_{id_2, t_2, p}$  such that  $loc(id_2, \{p\}, t_2) \notin C_2(e)$ , and the subsequent points in  $\nu$  (with index in  $[h+1, \dots, m]$ ) correspond to the variables  $v_{id_2, t_2, p}$  such that  $loc(id_2, \{p\}, t_2) \in C_2(e)$ . Let  $\nu_i$  be the sum of the values assigned by  $\sigma$  to the first  $i$  points in  $\nu$ , that is,  $\nu_i = \sum_{j=1}^i \sigma(v_{id_2, t_2, p_j})$  where  $p_i \in \nu$ . Observe that  $\nu_m = 1$  since equality (2) of Definition 15 holds for  $BC(\mathcal{K}^{\alpha+1})$ , and that  $\nu_h = 1 - \sum_{loc(id_2, \{p\}, t_2) \in C_2(e)} \sigma(v_{id,t,p})$ . The fact that inequality (4) of Definition 15 holds for  $BC(\mathcal{K}^{\alpha+1})$  entails that,

$$1 - \nu_h + \pi_k = \sum_{loc(id, \{p\}, t) \in C_2(e)} \sigma(v_{id,t,p}) + \sum_{loc(id, \{p\}, t) \in C_1(e)} \sigma(v_{id,t,p}) \leq 1.$$

Therefore, we obtain that  $\pi_k \leq \nu_h$ . Intuitively, this means that it is possible to define a PDF over the worlds for  $\mathcal{K}^{\alpha+1}$  so that no world satisfies two st-atoms in  $C(e)$ , while keeping satisfied all the PST atoms and std-formulas satisfied by  $M^\alpha$ . This is formally described below.

We define a model  $M^{\alpha+1}$  for  $\mathcal{K}^{\alpha+1}$  as follows. Let  $U = \langle u_1, u_2, \dots, u_z \rangle$  be the sequence consisting of the values in  $\{\pi_1, \dots, \pi_n\} \cup \{\nu_1, \dots, \nu_m\}$  ordered in ascending order. We define the following  $z$  non-zero probability worlds for  $\mathcal{K}^{\alpha+1}$ . For each  $u_i \in U$ , let  $w_i^*$  be a world such that:

- $\forall id \in ID \setminus \{id_2\}, \forall t \in T \setminus \{t_2\}, w_i^*(id, t) = w_j(id, t)$ , where  $w_j \in \mathcal{W}(\mathcal{K}^\alpha)$  and  $j$  is the smallest subscript such that  $\pi_j \geq u_i$ .
- $w_i^*(id_2, t_2) = p_j$ , where  $p_j \in S$  and  $j$  is the smallest subscript such that  $\nu_j \geq u_i$ .

We define  $M^{\alpha+1}$  as the PDF over  $\mathcal{W}(\mathcal{K}^{\alpha+1})$  such that (i)  $M^{\alpha+1}(w_1^*) = v_1$ , (ii)  $M^{\alpha+1}(w_i^*) = v_i - v_{i-1}$  for each  $i \in [1..z]$ , and (iii) for all the other worlds  $w \in \mathcal{W}(\mathcal{K}^{\alpha+1})$ ,  $M^{\alpha+1}(w) = 0$ .

We now show that  $M^{\alpha+1}$  is a model for  $\mathcal{K}^{\alpha+1}$ . First, observe that  $M^{\alpha+1}$  is a model for  $\mathcal{K}^\alpha$ . In fact, for each  $id \in ID \setminus \{id_2\}, id \in T \setminus \{t_2\}$ , and  $p$  in *Space*, the probability that object  $id$  is at time value  $t$  at point  $p$  according to  $M^{\alpha+1}$  is equal to the probability that  $id$  is at time value  $t$  at point  $p$  according to  $M^\alpha$ , that is

$$\forall id \in ID \setminus \{id_2\}, t \in T \setminus \{t_2\}, p \in Space, \quad \sum_{\substack{w \in \mathcal{W}(\mathcal{K}^{\alpha+1}), \\ w(id, t) = p}} M^{\alpha+1}(w) = \sum_{\substack{w \in \mathcal{W}(\mathcal{K}^\alpha), \\ w(id, t) = p}} M^\alpha(w).$$

Moreover, for the  $\langle id_2, t_2 \rangle$  pair,

$$\forall p \in Space, \quad \sum_{\substack{w \in \mathcal{W}(\mathcal{K}^{\alpha+1}), \\ w(id, t) = p}} M^{\alpha+1}(w) = \sigma(v_{id_2, t_2, p}) = \sum_{\substack{w \in \mathcal{W}(\mathcal{K}^\alpha), \\ w(id, t) = p}} M^\alpha(w).$$

This ensures that, for each PST atom  $a = loc(id, r, t)[\ell, u] \in \mathcal{A}$ , the sum of the probabilities assigned by  $M^{\alpha+1}$  to the worlds satisfying the st-atom  $loc(id, r, t)$  belong to the probability interval  $[\ell, u]$  specified by  $a$ , since the same hold for  $M^\alpha$ . Moreover, none of the worlds which is assigned a probability equal to zero by  $M^\alpha$  is assigned a probability greater than zero by  $M^{\alpha+1}$ , as  $\pi$  only consists of worlds which are assigned a non-zero probability by  $M^\alpha$ . Therefore,  $M^{\alpha+1}$  is a model for  $\mathcal{K}^\alpha$ .

Finally, to show that  $M^{\alpha+1}$  is a model for  $\mathcal{K}^{\alpha+1}$ , it suffices to observe that  $\pi_k \leq \nu_h$  entails that there is no world  $w_i^*$  for  $\mathcal{K}^{\alpha+1}$  (with  $i \in [1..z]$ ) which is assigned a probability greater than zero and which satisfies two atoms in  $C(e)$ . This completes the proof.  $\square$

The following example shows the usage of the result of Theorem 7 to check the consistency of a PST KB.

**Example 20.** Continuing Example 18, assume that  $\mathcal{K}$  is a PST KB where  $\mathcal{F}$  is the set of std-formulas given in Example 18 and  $\mathcal{A}$  consist of the following set PST atoms:

$loc(id_1, \{p_3, p_4\}, 0)[.7, 1]$ ,  $loc(id_1, \{p_1\}, 1)[.2, .5]$ ,  $loc(id_1, \{p_1\}, 2)[.3, .8]$ ,  $loc(id_2, \{p_2\}, t)[.7, 1]$  for  $t \in [0..2]$ . It can be checked that for this PST KB the corresponding set of linear inequalities  $BC(\mathcal{K})$  has a feasible solution. Thus, since the auxiliary graph  $\mathcal{AG}(\mathcal{K})$  shown in Figure 2(b) is acyclic, it follows that  $\mathcal{K}$  is consistent.

Consider now the PST KB  $\mathcal{K}'$  obtained from  $\mathcal{K}$  by replacing the atom  $loc(id_1, \{p_1\}, 2)[.3, .8]$  with  $loc(id_1, \{p_1\}, 2)[.8, .8]$ , where the lower probability is .8 instead of .3. In this case,  $BC(\mathcal{K}')$  includes the inequalities  $0.7 \leq v_{id_1, 0, p_3} + v_{id_1, 0, p_4} \leq 1$  and  $0.8 \leq v_{id_1, 2, p_1} \leq 0.8$  (due to PST atoms), and  $v_{id_1, 0, p_3} + v_{id_1, 0, p_4} + v_{id_1, 1, p_1} \leq 1$  (due to the edge  $\langle (id_1, 0), (id_1, 2) \rangle$  of  $\mathcal{AG}(\mathcal{K})$ ). Clearly, there is no feasible solution of  $BC(\mathcal{K}')$ , and thus  $\mathcal{K}'$  is not consistent.  $\square$

We note the fact that the acyclicity of  $\mathcal{AG}$  is not used in the left-to-right proof of Theorem 6. Therefore, whether or not  $\mathcal{AG}(\mathcal{F})$  is acyclic, for PST KBs where the std-formulas are binary, the necessary condition stated in the following theorem can be used for checking consistency.

**Theorem 7.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a PST KB where  $\mathcal{F}$  consists of binary std-formulas only. If there is no feasible solution of  $BC(\mathcal{K})$ , then  $\mathcal{K}$  is not consistent.*

The condition that  $BC(\mathcal{K})$  has no feasible solution is stronger than that  $NC(\mathcal{K})$  has no feasible solution, used in Theorem 3 for general std-formulas, since it can be easily checked that every solution of  $BC(\mathcal{K})$  is a solution of  $NC(\mathcal{K})$ , but the converse does not hold in general. Thus, there are PST KBs  $\mathcal{K}$ , such that using Theorem 3 we cannot conclude anything about the consistency of  $\mathcal{K}$ , as  $NC(\mathcal{K})$  has a feasible solution, but we can show that  $\mathcal{K}$  is inconsistent by using the result of Theorem 7.

**Example 21.** We have already observed that the fact that the PST KB  $\mathcal{K}$  of Example 16 is not consistent cannot be concluded by checking whether there is a feasible solution of  $NC(\mathcal{K})$ . However, it is easy to see that  $BC(\mathcal{K})$  has no feasible solution in this case, and thus we can conclude that the PST KB of Example 16 is not consistent by applying the result of Theorem 7.  $\square$

Although the necessary condition in Theorem 6 can be still used to check the consistency when the auxiliary graph is cyclic (see Theorem 7), the sufficient condition of Theorem 6 entails that a PST KB  $\mathcal{K}$  is consistent only if the corresponding set  $BC(\mathcal{K})$  of linear inequalities has a feasible solution and the auxiliary graph is acyclic. In the following, we identify a class of std-graphs for which the consistency checking problem remains tractable even if the std-graph is clique-cyclic (and thus the auxiliary graph is cyclic).

### 3.5.2 SIMPLE CLIQUE-CYCLIC STD-GRAPHS

In this section we provide a set of linear inequalities for checking the consistency of a PST KB whose std-graph satisfies the following property.

**Definition 16** (Simple clique-cyclic std-graph). *Std-graph  $\mathcal{G}$  is said to be simple clique-cyclic iff its auxiliary graph  $\mathcal{AG}$  is a simple graph<sup>11</sup> and each cyclic connected component of  $\mathcal{G}$  contains a single maximal clique that is not an  $\langle id, t \rangle$  clique.*

The following is an example of a simple clique-cyclic std-graph. We provide both the std-graph and its auxiliary graph.

**Example 22.** For the PST KB  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  of Example 15, we obtain the std-graph  $\mathcal{G}$  shown in Figure 3(a). The auxiliary graph  $\mathcal{AG}$  is shown in Figure 3(b).  $\mathcal{G}$  consists of a single connected component where there is one maximal clique  $Cl$  that is not an  $\langle id, t \rangle$  clique, namely  $Cl = \{loc(id, \{p_0\}, 0), loc(id, \{p_0\}, 1), loc(id, \{p_0\}, 2)\}$ . It is easy to see that  $\mathcal{G}$  is simple clique-cyclic.  $\square$

The following theorem states that checking the consistency of a PST KB whose std-graph is simple clique-cyclic can be accomplished by checking whether there is a feasible solution of a set of linear inequalities.

We start by defining this new linear system:  $CL(\mathcal{K})$ .

11. There are no loops and no multiple edges between vertices.

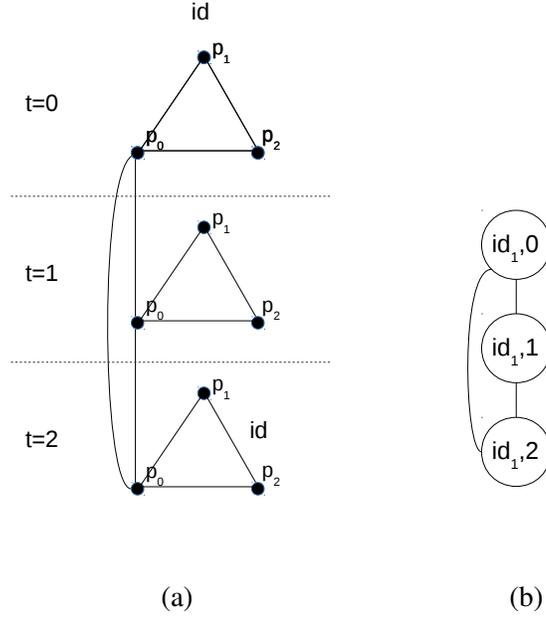


Figure 3: (a) Std-graph and (b) auxiliary graph for the PST KB of Example 15.

**Definition 17** ( $CL(\mathcal{K})$ ). Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a PST KB such that  $\mathcal{F}$  consists of binary std-formulas only.  $CL(\mathcal{K})$  consists of the following (in)equalities:

$$(1) \forall loc(id, r, t)[\ell, u] \in \mathcal{A}: \ell \leq \sum_{p \in r} v_{id,t,p} \leq u;$$

$$(2) \forall id \in ID, t \in T: \sum_{p \in Space} v_{id,t,p} = 1;$$

$$(3) \forall p \in Space, id \in ID, t \in T: v_{id,t,p} \geq 0;$$

(4) (a) For each acyclic connected component of  $\mathcal{G}$  (and hence  $\mathcal{AG}$ ), for each edge  $e$  of  $\mathcal{AG}$ :

$$\sum_{loc(id, \{p\}, t) \in C(e)} v_{id,t,p} \leq 1.$$

(b) For each cyclic connected component of  $\mathcal{G}$  (and hence  $\mathcal{AG}$ ), for the maximal clique  $Cl$  that is not an  $\langle id, t \rangle$  clique:

$$\sum_{loc(id, \{p\}, t) \in Cl} v_{id,t,p} \leq 1.$$

The intuition behind the result stated in the following theorem is similar to that given for the tractable case of Theorem 6. Indeed, if  $\mathcal{G}$  is simple clique-cyclic and  $CL(\mathcal{K})$  has a feasible solution, then the events corresponding to the st-atoms in each edge  $e$  of an acyclic connected component  $\mathcal{AG}$ , as well as the events corresponding to the st-atoms in each maximal clique of a connected component of  $\mathcal{AG}$ , can be arranged in the probability space avoiding overlaps. This basically suffices to define a model for  $\mathcal{K}$ .

**Theorem 8.** Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a PST KB where  $\mathcal{F}$  consists of binary std-formulas only. If  $\mathcal{G}$  is simple clique-cyclic, then  $\mathcal{K}$  is consistent iff there is a feasible solution of  $CL(\mathcal{K})$ .

*Proof.* As explained earlier we will deal only with a single cyclic connected component of  $\mathcal{G}$ .

( $\Rightarrow$ ) The proof follows by reasoning as in the left-to-right proof of Theorem 6.

( $\Leftarrow$ ) Let  $\sigma$  be a solution for  $CL(\mathcal{K})$ . Obviously,  $\sigma$  is also a solution for  $CC(\mathcal{K}')$  where  $\mathcal{K}' = \langle \mathcal{A}, \emptyset \rangle$ , and there is a model  $M'$  for  $\mathcal{K}'$  such that  $\sigma(v_{id,t,p}) = \sum_{w \in \mathcal{W}(\mathcal{K}'), w \models loc(id, \{p\}, t)} M'(w)$  and  $\sum_{loc(id, \{p\}, t) \in Cl} \sigma(v_{id,t,p}) \leq 1$ . Given this, we show that starting from  $\sigma$  we can define a new model  $M$  for  $\mathcal{K}'$  which is also a model for  $\mathcal{K}$ .

For each  $\langle id, t \rangle$  pair, we define  $S(id, t)$  as the subset of  $Space$  consisting of points  $p$  such that  $v_{id,t,p}$  is assigned a value greater than zero by  $\sigma$ , i.e.,  $S(id, t) = \{p \mid p \in Space, \sigma(v_{id,t,p}) > 0\}$ . We distinguish between the following two sets of  $\langle id, t \rangle$  pairs:  $IDT_1$  (resp.  $IDT_2$ ) is the set of  $\langle id, t \rangle$  pairs for which there is not (resp. there is) a point  $p$  in  $S(id, t)$  such that  $loc(id, \{p\}, t) \in Cl$ . We separately consider the pairs in  $IDT_1$  and  $IDT_2$ , and for each  $\langle id, t \rangle$  pair belonging to one of these sets, define a sequence of points in  $S(id, t)$  along with a sequence of cumulative probability values that will be used to build the model  $M$ .

We start with the set  $IDT_1$ . Let  $\langle id_i, t_i \rangle$ , with  $i \in [1..|IDT_1|]$  be the  $i$ -th pair of  $IDT_1$  (after ordering the pairs in  $IDT_1$  according to any fixed order). Let  $\pi_i = \langle p_i^1, \dots, p_i^{m_i} \rangle$  be a permutation of the points in  $S(id_i, t_i)$ , and let  $\pi_i(k)$  be the sum of the values assigned by  $\sigma$  to the first  $k$  points in  $\pi_i$ , that is,  $\pi_i(k) = \sum_{j=1}^k \sigma(v_{id_i, t_i, p_i^j})$ . Thus,  $\langle \pi_i(1), \dots, \pi_i(m_i) \rangle$  is a sequence of (cumulative) probability values associated with the sequence of points  $\langle p_i^1, \dots, p_i^{m_i} \rangle$  from the  $i$ -th pair of  $IDT_1$ .

Now consider the pairs in  $IDT_2$ , i.e. those occurring in some st-atom in  $Cl$ . Denote with  $\langle id_i, t_i \rangle$ , with  $i \in [1..|IDT_2|]$ , the  $i$ -th pair of  $IDT_2$  (according to any fixed order). For the first pair in  $IDT_2$ , let  $\nu_1 = \langle p_1^1, \dots, p_1^{h_1}, p_1^{h_1+1}, \dots, p_1^{m_1} \rangle$  be a permutation of the points in  $S(id_1, t_1)$  such that the first  $h_1$  points  $p$  correspond to the variables  $v_{id_1, t_1, p}$  such that  $loc(id_1, \{p\}, t_1) \in Cl$  (consequently, the other points  $p \in \{p_1^{h_1+1}, \dots, p_1^{m_1}\}$  correspond to variables  $v_{id_1, t_1, p}$  such that  $loc(id_1, \{p\}, t_1) \notin Cl$ ).

For each  $\langle id_i, t_i \rangle$  with  $i \in [2..|IDT_2|]$ , let  $\nu_i = \langle p_i^1, \dots, p_i^{g_i}, p_i^{g_i+1}, \dots, p_i^{g_i+h_i}, p_i^{g_i+h_i+1}, \dots, p_i^{m_i} \rangle$  be a sequence of the points in  $S(id_i, t_i)$  such that (i) each point  $p$  in  $\{p_i^1, \dots, p_i^{g_i}, p_i^{g_i+h_i+1}, \dots, p_i^{m_i}\}$  (resp.  $p$  in  $\{p_i^{g_i+1}, \dots, p_i^{g_i+h_i}\}$ ) corresponds to a variable  $v_{id_i, t_i, p}$  such that  $loc(id_i, \{p\}, t_i) \notin Cl$  (resp.  $loc(id_i, \{p\}, t_i) \in Cl$ ); and (ii) all the points in the sequence are distinct except for  $p_i^{g_i}$  and  $p_i^{g_i+h_i+1}$  that may refer to the same point in  $S(id_i, t_i)$ . Denoting as  $\nu_i(k)$  the sum of the values assigned by  $\sigma$  to the first  $k$  points in  $\nu_i$  (i.e.,  $\nu_i(k) = \sum_{j=1}^k \sigma(v_{id_i, t_i, p_i^j})$ ), we choose the sequences  $\nu_2, \dots, \nu_{|IDT_2|}$  such that, for each  $i \in [2..|IDT_2|]$ ,  $\nu_i(g_i) = \nu_{i-1}(g_{i-1} + h_{i-1})$  (assume that  $g_1 = 0$ ). Note that we can make such a choice since the following holds:

$$\sum_{i=1}^{|IDT_2|} \sum_{k=1}^{h_i} \sigma(v_{id_i, t_i, p_i^k}) = \sum_{i=1}^{|IDT_2|} \nu_i(g_i + h_i) - \nu_i(g_i) = \sum_{i=1}^{|IDT_2|} \sum_{loc(id_i, \{p\}, t_i) \in Cl} \sigma(v_{id_i, t_i, p}) \leq 1.$$

The way the sequences  $\langle \nu_1, \dots, \nu_{|IDT_2|} \rangle$  and  $\langle \nu_i(1), \dots, \nu_i(m_i) \rangle$  are defined will allow us to build a model  $M$  for  $\mathcal{K}'$  such that no world which is assigned a probability greater than zero satisfies two distinct st-atoms in  $Cl$ . This will ensure that  $M$  is a model for  $\mathcal{K}$  too.

We now define such a model  $M$ . Let  $V = \langle v_1, v_2, \dots, v_z \rangle$  be the sequence consisting of the values in  $\bigcup_{i \in IDT_1} \{\pi_i(1), \dots, \pi_i(m_i)\} \cup \bigcup_{i \in IDT_2} \{\nu_i(1), \dots, \nu_i(m_i)\}$  ordered in ascending order. We define the following  $z$  non-zero probability worlds for  $\mathcal{K}'$  (note that a world for  $\mathcal{K}'$  is also a world for  $\mathcal{K}$ ). For each  $v_j \in V$ , let  $w_j$  be a world such that:

- for each  $\langle id_i, t_i \rangle \in IDT_1$ , we define  $w_j(id_i, t_i) = p_i^k$ , where  $p_i^k$  is the left-most value in  $\langle \pi_i(1), \dots, \pi_i(m_i) \rangle$  such that  $\pi_i(k) \geq v_j$ .
- for each  $\langle id_i, t_i \rangle \in IDT_2$ , we define  $w_j(id_i, t_i) = p_i^k$ , where  $p_i^k$  is the left-most value in  $\langle \nu_i(1), \dots, \nu_i(m_i) \rangle$  such that  $\nu_i(k) \geq v_j$ .

Finally, we define  $M$  as the PDF over  $\mathcal{W}(\mathcal{K}')$  such that (i)  $M(w_1) = v_1$ , (ii)  $M(w_j) = v_j - v_{j-1}$  for each  $j \in [1..z]$ , and (iii)  $M(w) = 0$  for all the other worlds  $w \in \mathcal{W}(\mathcal{K}')$ . It is easy to check that  $M$  is a model for  $\mathcal{K}'$  that corresponds to the same solution  $\sigma$  for  $CC(\mathcal{K}')$ , that is,  $M$  is such that  $\sigma(v_{id,t,p}) = \sum_{w \in \mathcal{W}(\mathcal{K}'), w \models loc(id, \{p\}, t)} M(w)$ . Moreover, as the construction shown above ensures that no non-zero probability world for  $\mathcal{K}'$  satisfies two distinct st-atoms in  $Cl$ , it follows that  $M$  is a model for  $\mathcal{K}$ .  $\square$

Theorem 8 can be used, for instance, to decide if the PST KB of Example 15 is consistent.

**Example 23.** For the PST KB  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  of Example 15,  $CL(\mathcal{K})$  is the linear system obtained from  $CC(\langle \mathcal{A}, \emptyset \rangle)$  by augmenting it with the inequality  $v_{id,0,p_0} + v_{id,1,p_0} + v_{id,2,p_0} \leq 1$ . Since  $CL(\mathcal{K})$  also contains the inequality  $0.5 \leq v_{id,t,p_0}$  for each  $t \in [0..2]$  (due to the presence of the PST atom  $loc(id, \{p_0\}, t)[0.5, 1]$  in  $\mathcal{A}$ ), it follows that  $CL(\mathcal{K})$  has no feasible solution, from which we conclude that  $\mathcal{K}$  is not consistent (here  $\mathcal{G}$  is the clique-cyclic std-graph shown in Figure 3(a)).  $\square$

We conclude the section with the following theorem stating that checking consistency for the PST KBs identified in Section 3.5 is tractable.

**Theorem 9.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a PST KB where  $\mathcal{F}$  consists of binary std-formulas only. If  $\mathcal{G}$  is clique-acyclic or simple clique-cyclic, then checking the consistency of  $\mathcal{K}$  is in PTIME.*

*Proof.* If  $\mathcal{G}$  consists of one connected component, the statement follows from the facts that building  $\mathcal{G}$ , as well as checking whether  $\mathcal{G}$  is acyclic (resp. simple clique-cyclic), can be accomplished in polynomial time w.r.t. the size of  $\mathcal{K}$  and checking whether there is a feasible solution of  $BC(\mathcal{K})$  (resp.,  $CL(\mathcal{K})$ ) is polynomial too. If  $\mathcal{G}$  consists of more than one connected component, the statement follows from the facts that finding all connected components of  $\mathcal{G}$  can be done in polynomial time, and, using Proposition 3, we only need to check the consistency of the PST KBs corresponding to each of the connected components in order to decide the consistency of the whole PST KB.  $\square$

## 4. Querying PST Knowledge Bases

This section investigates the complexity of checking answer to queries. Section 4.1 contains the basic definitions. Section 4.2 contains the major result about the complexity of checking answers to queries. Finally, Section 4.3 gives sufficient and necessary conditions for answering queries and tractable cases.

### 4.1 Optimistic and Cautious Answers

The problem of querying SPOT data has been investigated by Parker et al. (2007a, 2009), Parisi, Parker, Grant, and Subrahmanian (2010), and Grant, Molinaro, and Parisi (2013) in more specific

frameworks corresponding to PST KBs (of the form  $\langle \mathcal{A}, \emptyset \rangle$ ), where only PST atoms were considered. In this section, we address the problem of answering *selection queries* in general (consistent) PST KBs. These kinds of queries were considered by Parker et al. (2007a, 2009), and Parisi et al. (2010), while Grant et al. (2013) focused on count queries.

A selection query is an expression of the form  $(?id, q, ?t, [\ell, u])$ , where  $q$  is a region and  $[\ell, u]$  is a probability interval. Intuitively, a selection query says: “Given a region  $q$  and a probability interval  $[\ell, u]$ , find all objects  $id$  and times  $t$  such that  $id$  is inside  $q$  at time  $t$  with a probability in the interval  $[\ell, u]$ .” There are two semantics for interpreting this statement, leading to two types of answers to selection queries. *Optimistic answers* are objects and time values that *may* be in the query region with probability in the specified interval, whereas *cautious answers* consist only of those objects and time values that are *guaranteed* to be in that region with probability in the given interval. Thus, the cautious answers are a subset of the optimistic ones.

The following definition extends the original definition of optimistic and cautious selection query answers to the general case of consistent PST knowledge bases.

**Definition 18** (Optimistic/Cautious Query Answers). *Let  $\mathcal{K}$  be a consistent PST KB, and  $Q = (?id, q, ?t, [\ell, u])$  a selection query. Then,*

- $\langle id, t \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$  iff there exists a model  $M \in \mathbf{M}(\mathcal{K})$  s.t.
 
$$\sum_{w \models loc(id, q, t)} M(w) \in [\ell, u].$$
- $\langle id, t \rangle$  is a cautious answer to  $Q$  w.r.t.  $\mathcal{K}$  iff for every model  $M \in \mathbf{M}(\mathcal{K})$ ,
 
$$\sum_{w \models loc(id, q, t)} M(w) \in [\ell, u].$$

**Example 24.** Let  $q_1 = \{(7, 3), (7, 4)\}$  ( $q_1$  overlaps with region  $c$ , see Figure 1). Model  $M$  of Example 9 entails that  $\langle id_1, 9 \rangle$  is an optimistic answer to  $Q = (?id, q_1, ?t, [.7, 1])$ , as  $w_1(id_1, 9) = w_2(id_1, 9) = w_3(id_1, 9) = (7, 4) \in q_1$  and  $M(w_1) + M(w_2) + M(w_3) = 1 \in [.7, 1]$ .

Let  $q_2$  be any region including region  $c$ .  $\langle id_1, 9 \rangle$  is a cautious answer to  $Q' = (?id, q_2, ?t, [.7, 1])$ , as according to any model for  $\mathcal{K}_{ex}$ ,  $id_1$  must be in region  $c$  (and thus in  $q_2$ ) at time 9 with probability in  $[.9, 1]$  (due to  $loc(id_1, c, 9)[.9, 1] \in \mathcal{A}_{ex}$ ). Clearly,  $\langle id_1, 9 \rangle$  is also an optimistic answer to  $Q'$ .  $\square$

## 4.2 Complexity of Query Answering

The following theorem shows how consistency checking can be used to answer selection queries under both the optimistic and the cautious semantics.

**Theorem 10.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a consistent PST KB, and  $Q = (?id, q, ?t, [\ell, u])$ . Then,*

- 1)  $\langle id, t \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$  iff  $\langle \mathcal{A} \cup \{loc(id, q, t)[\ell, u]\}, \mathcal{F} \rangle$  is consistent.
- 2)  $\langle id, t \rangle$  is a cautious answer to  $Q$  w.r.t.  $\mathcal{K}$  iff both  $\langle \mathcal{A} \cup \{loc(id, q, t)[0, \ell - \epsilon]\}, \mathcal{F} \rangle$  and  $\langle \mathcal{A} \cup \{loc(id, q, t)[u + \epsilon, 1]\}, \mathcal{F} \rangle$  are not consistent, where  $\epsilon$  is the smallest rational number with precision  $m^3$ , where  $m = |\mathcal{A}| + |\mathcal{F}| \cdot |ID| \cdot |T| \cdot |Space| + |ID| \cdot |T| + |ID| \cdot |T| \cdot |Space|$ .

*Proof.* Statement 1) easily follows from Definitions 6 and 18.

Statement 2) follows from the fact that  $\langle id, t \rangle$  is a cautious answer to  $Q$  w.r.t.  $\mathcal{K}$  iff i)  $\langle id, t \rangle$  is not an optimistic answer to  $Q' = (?id, q, ?t, [0, \ell - \epsilon])$  w.r.t.  $\mathcal{K}$ , and ii)  $\langle id, t \rangle$  is not an optimistic

answer to  $Q'' = (?id, q, ?t, [u + \epsilon, 1])$  w.r.t.  $\mathcal{K}$ , where  $\epsilon$  is a “small” non-zero constant whose size is polynomial w.r.t. the size of  $\mathcal{K}$ . To show the existence of  $\epsilon$ , we observe that (i) checking the consistency of  $\mathcal{K}$  can be reduced to deciding the PSAT instance  $\phi_{\mathcal{K}}$  consisting of  $m$  clauses, where  $m = |\mathcal{A}| + |\mathcal{F}| \cdot |ID| \cdot |T| \cdot |Space| + |ID| \cdot |T| + |ID| \cdot |T| \cdot |Space|$ , as shown in the membership proof of Theorem 1; and (ii) if a satisfying probability distribution for an instance of PSAT with  $m$  clauses exists, then there is one with at most  $m + 1$  non-zero probabilities, and with entries consisting of rational numbers with precision  $O(m^2)$  (Georgakopoulos et al., 1988; Papadimitriou & Steiglitz, 1982); Then, choosing  $\epsilon$  equal to the smallest rational number with precision  $m^3$  suffices to obtain a sufficiently small non-zero constant, whose size is polynomial w.r.t. the size of  $\mathcal{K}$ , such that  $\langle \mathcal{A} \cup \{loc(id, q, t)[0, \ell - \epsilon]\}, \mathcal{F} \rangle$  (resp.,  $\langle \mathcal{A} \cup \{loc(id, q, t)[u + \epsilon, 1]\}, \mathcal{F} \rangle$ ) is consistent iff  $\langle \mathcal{A} \cup \{loc(id, q, t)[0, \ell - \nu]\}, \mathcal{F} \rangle$  (resp.,  $\langle \mathcal{A} \cup \{loc(id, q, t)[u + \nu, 1]\}, \mathcal{F} \rangle$ ) is consistent, where  $\nu$  is an infinitely small value greater than zero.  $\square$

The following corollary follows from Theorems 1 and 10.

**Corollary 1.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a consistent PST KB. Given a query  $Q = (?id, q, ?t, [\ell, u])$ , then*

- 1) *deciding whether  $\langle id, t \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$  is in NP.*
- 2) *deciding whether  $\langle id, t \rangle$  is a cautious answer to  $Q$  w.r.t.  $\mathcal{K}$  is coNP.*

The problem of deciding whether a pair  $\langle id, t \rangle$  is an optimistic/cautious answer to a selection query  $Q$  can be solved in polynomial time for PST KBs whose  $\mathcal{F}$  component is empty (Parker et al., 2007a, 2009; Parisi et al., 2010). However, in the presence of integrity constraints this problem becomes hard.

**Theorem 11.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a consistent PST KB. Given a query  $Q = (?id, q, ?t, [\ell, u])$ , then*

- 1) *deciding whether  $\langle id, t \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$  is NP-hard.*
- 2) *deciding whether  $\langle id, t \rangle$  is a cautious answer to  $Q$  w.r.t.  $\mathcal{K}$  is coNP-hard.*

*Proof.* We first prove item 1). We show a reduction to our problem from the NP-hard Hamiltonian path problem (Papadimitriou, 1994). Given a directed graph  $G = \langle V, E \rangle$ , where  $V = \{v_0, \dots, v_k\}$  is the set of  $k + 1$  vertexes of  $G$  and  $E$  is the set of its edges, we construct an instance of our problem as follows. Let  $ID = \{id\}$ ,  $Space = V \cup \{p_0, \dots, p_{k-1}\} \cup \{p_T, p_F\}$ , and  $T = [0, \dots, 2k + 1]$ .  $\mathcal{K}$  is the pair  $\langle \mathcal{A}, \mathcal{F} \rangle$  such that  $\mathcal{A}$  consist of the PST atom  $loc(id, v_0, 0)[1, 1]$  and  $\mathcal{F}$  consists of the following formulas:

- $\forall i \in [0..k - 1], f_1^i = \forall Z_1, Z_2 \neg [loc(id, \{v_i\}, Z_1) \wedge loc(id, Space \setminus V', Z_2) \wedge Z_2 = Z_1 + 1]$  where  $V' = \{v_j | (v_i, v_j) \in E\} \cup \{p_i, p_T, p_F\}$ . That is, the locations  $id$  can reach starting from  $v_i$  in one time step are only the locations  $v_j$  s.t.  $(v_i, v_j) \in E$  or any point in  $\{p_i, p_T, p_F\}$ .
- $f_1^k = \forall Z_1, Z_2 \neg [loc(id, \{v_k\}, Z_1) \wedge loc(id, Space \setminus V'', Z_2) \wedge Z_2 = Z_1 + 1]$  where  $V'' = \{v_j | (v_k, v_j) \in E\} \cup \{p_T, p_F\}$ . This formula says that the locations  $id$  can reach starting from  $v_k$  in one time step are only locations  $v_j$  such that  $(v_k, v_j) \in E$  or those in  $\{p_T, p_F\}$ .
- $\forall i \in [0..k - 1], f_2^i = \forall Z_1, Z_2 \neg [loc(id, \{p_i\}, Z_1) \wedge loc(id, Space \setminus \{v_{i+1}\}, Z_2) \wedge Z_2 = Z_1 + 1]$ , saying that the only location  $id$  can reach starting from  $p_i$  in one time step is  $v_{i+1}$ .

- $f_3 = \forall Y_1, Z_1, Z_2 \neg[loc(id, Y_1, Z_1) \wedge loc(id, Y_1, Z_2) \wedge Z_1 \neq Z_2 \wedge Y_1 \neq \{p_T\} \wedge Y_1 \neq \{p_F\}]$ , saying that  $id$  can not be on the same location at distinct time values, if the location is different from both  $p_T$  and  $p_F$ .
- $f_4 = \forall Z_1, Z_2 \neg[loc(id, \{p_0, \dots, p_{k-1}\}, Z_1) \wedge loc(id, \{p_T\}, Z_2)]$ , saying that  $id$  can not be at  $p_T$  if it was/is/will be at a location in  $\{p_0, \dots, p_{k-1}\}$ .
- $f_5 = \forall Y_2, Z_1, Z_2 \neg[loc(id, \{p_T\}, Z_1) \wedge loc(id, Y_2, Z_2) \wedge Y_2 \neq \{p_T\} \wedge Z_1 < Z_2]$ , saying that  $id$  can not go away from  $p_T$  after reaching this location.
- $f_6 = \forall Y_2, Z_1, Z_2 \neg[loc(id, \{p_F\}, Z_1) \wedge loc(id, Y_2, Z_2) \wedge Y_2 \neq \{p_F\} \wedge Z_1 < Z_2]$ , saying that  $id$  can not go away from  $p_F$  after reaching this location.

Finally, let  $Q = (?id, \{p_T\}, ?t, [1, 1])$ , and let the pair to be checked as an the optimistic answer of  $Q$  be  $\langle id, 2k + 1 \rangle$ .

First of all, observe that  $\mathbf{M}(\mathcal{K}) \neq \emptyset$ . In fact, consider the world  $w \in \mathcal{W}(\mathcal{K})$  such that  $w(id, 0) = v_0$ ,  $w(id, 1) = p_0$ ,  $w(id, 2) = v_1$ ,  $w(id, 3) = p_1, \dots, w(id, t) = v_{t/2}$ ,  $w(id, t + 1) = p_{t/2+1}$ , (in both cases  $t$  is even)  $\dots, w(id, 2k-2) = v_{k-1}$ ,  $w(id, 2k-1) = p_k$ ,  $w(id, 2k) = v_k$ ,  $w(id, 2k+1) = p_F$ . It is easy to see that  $w \models \mathcal{F}$ . In fact, for each  $t \in [0..2k]$ ,  $id$  moves towards reachable locations in one time point: if  $id$  is at  $v_i$ , with  $i \in [0..k - 1]$ , it moves to  $p_i$  (that is,  $f_1^i$  is satisfied by  $w$ ); if  $id$  is at  $p_i$ , with  $i \in [0..k - 1]$ , it moves to  $v_{i+1}$  (that is,  $f_2^i$  is satisfied by  $w$ ); and when is at  $v_k$  it moves to  $p_F$  (that is,  $f_1^k$  is satisfied by  $w$ ). Moreover,  $id$  is not placed by  $w$  on the same location at different times (that is,  $f_3$  is satisfied), neither at both a location in  $\{p_1, \dots, p_{k-1}\}$  and at  $p_T$  (that is,  $f_4$  is satisfied too). Moreover,  $id$  does not move from  $p_F$  after reaching this location (that is,  $f_6$  is satisfied), and does not reach location  $p_T$  (thus, also  $f_5$  is satisfied by  $w$ ). Now, consider the PDF  $I$  over  $\mathcal{W}(\mathcal{K})$  assigning probability equal to 1 to  $w$ , and probability equal to 0 to every other world in  $\mathcal{W}(\mathcal{K})$ . Clearly,  $I$  is a model for  $\mathcal{K}$ : it assigns non-zero probability only to worlds satisfying  $\mathcal{F}$ , and it ensures that atom  $loc(id, v_0, 0)[1, 1]$  is satisfied by  $w$  which is assigned probability equal to 1.

We now show that  $\langle id, 2k + 1 \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$  iff there is a Hamiltonian path in  $G$ .

( $\Rightarrow$ ) As  $\langle id, 2k + 1 \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$ , there is a model  $M \in \mathbf{M}(\mathcal{K})$  s.t.  $\sum_{w|w \models loc(id, p_T, 2k+1)} M(w) = 1$ . Let  $W \subseteq \mathcal{W}(\mathcal{K})$  be the set of worlds  $w$  such that  $w(id, 2k + 1) = p_T$  and  $M(w) > 0$ . As  $\sum_{w \in W} M(w) = 1$  and  $M$  is a model (thus,  $\sum_{w|w \models loc(id, v_0, 0)} M(w) = 1$  due to  $loc(id, v_0, 0)[1, 1] \in \mathcal{A}$ ), it follows that every world  $w \in W$  is such that  $w(id, 0) = v_0$ . Moreover, since  $M$  is a model, and thus  $f_4$  is satisfied by every world  $w \in W$ , it holds that  $\forall t \in [0..2k + 1]$ ,  $w(id, t) \notin \{p_0, \dots, p_{k-1}\}$ , meaning that  $id$  is never placed at a location in  $\{p_0, \dots, p_{k-1}\}$  by  $w$ . Moreover, since for each  $w \in W$ ,  $w(id, 0) = v_0$ , and since  $f_1^i$  and  $f_2^i$  (with  $i \in [1..k]$ ), and  $f_3$  are satisfied by every world in  $W$ , it follows that (i) for each  $t \in [0, k]$ ,  $w(id, t) \in V$  (i.e.,  $id$  is placed on a vertex of  $G$  at each time point in  $[0, k]$ ), and (ii) for each  $t, t' \in [0, k]$ ,  $w(id, t) \neq w(id, t')$  ( $id$  is not placed on the same vertex of  $G$  at different times in  $[0, k]$ ). Given this, at time  $t = k + 1$ ,  $f_1^k$  and  $f_3$  entail that  $id$  is placed by any world in  $W$  at location  $p_T$  or  $p_F$ . As  $f_5$  (resp.,  $f_6$ ) entails that  $id$  can not go away from  $p_T$  (resp.,  $p_F$ ) after reaching this location, and since  $w(id, 2k + 1) = p_T$ , it follows that every world  $w \in W$  is such that for each  $t \in [k + 1, 2k + 1]$ ,  $w(id, t) = p_T$ . Hence, each  $w \in W$  is such that  $M(w) > 0$ ,  $w(id, 0) = v_0$ ,  $\forall t \in [1, k]$ ,  $w(id, t) \in V$ ,  $\forall t, t' \in [0, k]$ ,  $w(id, t) \neq w(id, t')$ ,  $\forall t \in [k + 1, 2k + 1]$ ,  $w(id, t) = p_T$ . Therefore,  $w(id, 0), w(id, 1), \dots, w(id, k)$  is a Hamiltonian path of  $G$ .

( $\Leftarrow$ ) Let  $\pi$  be a Hamiltonian path of  $G$ . We denote as  $\pi[i]$  (with  $i \in [0..k]$ ) the  $i$ -th vertex of  $\pi$ . W.l.o.g. we assume that the first vertex of  $\pi$  is  $v_0$ , that is,  $\pi[0] = v_0$ . We show that there is a model  $M \in \mathbf{M}(\mathcal{K})$  such that  $\sum_{w|w \models \text{loc}(id, p_T, 2k+1)} M(w) = 1$ , that is,  $\langle id, 2k+1 \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$ . Let  $M$  be a function over  $\mathcal{W}$  such that for all worlds  $w \in \mathcal{W}$ ,  $M(w) = 0$  except for the world  $w^*$  which is such that:  $w^*(id, 0) = \pi[0] = v_0$ ,  $\forall t \in [1, k]$ ,  $w^*(id, t) = \pi[t]$ ,  $\forall t \in [k+1, 2k+1]$ ,  $w^*(id, t) = p_T$ . It is easy to see that  $w^* \models \mathcal{F}$ . In fact, for each  $i \in [1..k]$ ,  $f_1^i$  as well as  $f_2^i$  are satisfied by  $w^*$ , since the fact that  $\pi$  is a path on  $G$  entails that  $\forall t \in [0, k-1]$ ,  $w^*(id, t) = v_i$  and  $w^*(id, t+1) = v_j$  only if edge  $(v_i, v_j)$  is an edge of  $G$  and  $w^*(id, k+1) = p_T$  which is reachable from any location in  $V$ . Moreover,  $f_3$  is satisfied by  $w^*$ , since the fact that  $\pi$  is a Hamiltonian path entails that  $w^*$  places  $id$  on different locations at different times, except for location  $p_T$ . Formula  $f_4$  is satisfied by  $w^*$ , since  $w^*$  does not place  $id$  at any location in  $\{p_0, \dots, p_{k-1}\}$ . Formula  $f_5$  is satisfied by  $w^*$ , since  $w^*$  does not place  $id$  on a location different from  $p_T$  after time point  $k+1$ , when  $id$  reaches  $p_T$ . Finally,  $f_6$  is satisfied by  $w^*$  as well, since  $w^*$  does not place  $id$  on location  $p_F$ . Since  $w^* \models \mathcal{F}$ , it can be assigned by  $M$  a probability different from 0. Let  $M(w^*) = 1$ . Moreover, as  $\sum_{w|w \models \text{loc}(id, v_0, 0)} M(w) = M(w^*) = 1$ , the condition required by  $\text{loc}(id, v_0, 0)[1, 1] \in \mathcal{A}$  holds too, proving that  $M$  is model for  $\mathcal{K}$ .

We now prove item 2). We show a reduction to our problem from the complement of the Hamiltonian path problem. Given a directed graph  $G$ , we construct an instance of our problem as in the proof of 1), except that  $Q = (?id, \{p_T\}, ?t, [0, 1 - \epsilon])$ , where  $\epsilon > 0$  is as stated in Theorem 10. As shown in the proof of 1),  $\mathcal{K}$  is consistent. Moreover, there is no Hamiltonian path in  $G$  iff the pair  $\langle id, 2k+1 \rangle$  is not an the optimistic answer to  $Q' = (?id, \{p_T\}, ?t, [1, 1])$ . Now observe that  $\langle id, 2k+1 \rangle$  is not an the optimistic answer to  $Q'$  iff  $\langle id, 2k+1 \rangle$  is a cautious answer to  $Q = (?id, \{p_T\}, ?t, [0, 1 - \nu])$  for  $\nu > 0$  sufficiently ‘‘small’’. By reasoning as in the proof of Theorem 10, we can choose  $\nu = \epsilon$ , which ensures that if there is an  $M \in \mathbf{M}(\mathcal{K})$  such that  $\sum_{w|w \models \text{loc}(id, p_T, t)} M(w) < 1$ , then  $\sum_{w|w \models \text{loc}(id, p_T, t)} M(w) \leq 1 - \epsilon$ .  $\square$

Our query answering problem is related to the decision version of the entailment problem in probabilistic logic (Nilsson, 1986; Georgakopoulos et al., 1988), which is as follows. Given a conjunction  $\phi$  of clauses, each of them associated with a probability, and an additional clause  $C$  with given lower and upper bounds on the admissible values for its probability, decide whether there is a probability distribution satisfying the probability of all the clauses in  $\phi$  and the lower and upper bounds for  $C$ . As in probabilistic logic, where the entailment problem can be reduced to PSAT (Nilsson, 1986; Georgakopoulos et al., 1988), we have shown that the query answering problem in PST KBs can be reduced to the consistency checking problem. However, there is an important difference between the query answering problem and the probabilistic entailment problem: we assume that the input PST KBs is consistent, while the entailment problem is defined for any  $\phi$  (even if  $\phi$  is not satisfiable). As a consequence, the coNP-hardness of the probabilistic entailment problem straightforwardly follows from the NP-hardness of PSAT ( $\phi$  is unsatisfiable iff any contradiction  $C$  with any lower and upper bounds is entailed by  $\phi$ ). That is, satisfiability is a source of complexity for the probabilistic entailment problem. However, in our setting, we are given a consistent PST KB and thus checking the consistency cannot be a source of complexity for the query answering problem. In fact, Theorem 11 can be viewed as strengthening the coNP-hardness of the probabilistic entailment problem, as it is proved in a more specific setting corresponding to the case where the input formula  $\phi$  is known to be satisfiable.

### 4.3 Sets of Linear Inequalities for Answering Queries

The following corollary of Theorems 2 and 10 states that the set of mixed-binary linear inequalities  $MBL(\cdot)$  introduced in Definition 8 can be also exploited for answering queries.

**Corollary 2.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a consistent PST KB, and  $Q = (?id, q, ?t, [\ell, u])$  a query. Then,*

- 1) *if there is a feasible solution of  $MBL(\langle \mathcal{A} \cup \{loc(id, q, t)[\ell, u]\}, \mathcal{F})$ , then  $\langle id, t \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$ .*
- 2) *if  $MBL(\langle \mathcal{A} \cup \{loc(id, q, t)[0, \ell - \epsilon]\}, \mathcal{F})$  or  $MBL(\langle \mathcal{A} \cup \{loc(id, q, t)[u + \epsilon, 1]\}, \mathcal{F})$  has a feasible solution, then  $\langle id, t \rangle$  is not a cautious answer to  $Q$  w.r.t.  $\mathcal{K}$  (where  $\epsilon$  is as given in Theorem 10).*

Similarly, from Theorems 3 and 10 the following corollary follows, stating that the linear system  $NC(\cdot)$  introduced in Definition 9 can be also used for answering queries.

**Corollary 3.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a consistent PST KB, and  $Q = (?id, q, ?t, [\ell, u])$  a query. Then,*

- 1) *if there is no feasible solution of  $NC(\langle \mathcal{A} \cup \{loc(id, q, t)[\ell, u]\}, \mathcal{F})$ , then  $\langle id, t \rangle$  is not an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$ .*
- 2) *if both  $NC(\langle \mathcal{A} \cup \{loc(id, q, t)[0, \ell - \epsilon]\}, \mathcal{F})$  and  $NC(\langle \mathcal{A} \cup \{loc(id, q, t)[u + \epsilon, 1]\}, \mathcal{F})$  have no feasible solution, then  $\langle id, t \rangle$  is a cautious answer to  $Q$  w.r.t.  $\mathcal{K}$  (where  $\epsilon$  is as given in Theorem 10).*

Moreover, we obtain the following corollary of Theorems 4 and 10.

**Corollary 4.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a consistent PST KB where  $\mathcal{F} = \{\neg[loc(id_1, \{p_1\}, t_1) \wedge loc(id_2, \{p_2\}, t_2) \wedge \dots \wedge loc(id_n, \{p_n\}, t_n)]\}$ , and  $Q = (?id, q, ?t, [\ell, u])$  be a query. Then,*

- 1)  *$\langle id, t \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$  iff there is a feasible solution of  $NC(\langle \mathcal{A} \cup \{loc(id, q, t)[\ell, u]\}, \mathcal{F})$ .*
- 2)  *$\langle id, t \rangle$  is a cautious answer to  $Q$  w.r.t.  $\mathcal{K}$  iff both  $NC(\langle \mathcal{A} \cup \{loc(id, q, t)[0, \ell - \epsilon]\}, \mathcal{F})$  and  $NC(\langle \mathcal{A} \cup \{loc(id, q, t)[u + \epsilon, 1]\}, \mathcal{F})$  have no feasible solution, where  $\epsilon$  is as given in Theorem 10.*

From Theorems 6 and 10 we obtain the following result.

**Corollary 5.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a consistent PST KB where  $\mathcal{F}$  consists of binary std-formulas only, and  $Q = (?id, q, ?t, [\ell, u])$  be a query. If  $\mathcal{AG}$  is acyclic, then*

- 1)  *$\langle id, t \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$  iff there is a feasible solution of  $BC(\langle \mathcal{A} \cup \{loc(id, q, t)[\ell, u]\}, \mathcal{F})$ .*
- 2)  *$\langle id, t \rangle$  is a cautious answer to  $Q$  w.r.t.  $\mathcal{K}$  iff both  $BC(\langle \mathcal{A} \cup \{loc(id, q, t)[0, \ell - \epsilon]\}, \mathcal{F})$  and  $BC(\langle \mathcal{A} \cup \{loc(id, q, t)[u + \epsilon, 1]\}, \mathcal{F})$  have no feasible solution, where  $\epsilon$  is as given in Theorem 10.*

Moreover, the results of Theorem 7, where the shape of the auxiliary graph is not considered, and Theorem 10 entail the following.

**Corollary 6.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a consistent PST KB where  $\mathcal{F}$  consists of binary std-formulas only, and  $Q = (?id, q, ?t, [\ell, u])$  be a query. Then*

- 1) *if there is no feasible solution of  $BC(\langle \mathcal{A} \cup \{loc(id, q, t)[\ell, u]\}, \mathcal{F})$ , then  $\langle id, t \rangle$  is not an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$*
- 2) *if both  $BC(\langle \mathcal{A} \cup \{loc(id, q, t)[0, \ell - \epsilon]\}, \mathcal{F})$  and  $BC(\langle \mathcal{A} \cup \{loc(id, q, t)[u + \epsilon, 1]\}, \mathcal{F})$  have no feasible solution, then  $\langle id, t \rangle$  is a cautious answer to  $Q$  w.r.t.  $\mathcal{K}$  (where  $\epsilon$  is as given in Theorem 10).*

Still, from Theorems 8 and 10 we obtain the following corollary.

**Corollary 7.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a consistent PST KB where  $\mathcal{F}$  consists of binary std-formulas only, and  $Q = (?id, q, ?t, [\ell, u])$  be a query. If  $\mathcal{AG}$  is simple-clique cyclic, then*

- 1)  *$\langle id, t \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}$  iff there is a feasible solution of  $CL(\langle \mathcal{A} \cup \{loc(id, q, t)[\ell, u]\}, \mathcal{F})$ .*
- 2)  *$\langle id, t \rangle$  is a cautious answer to  $Q$  w.r.t.  $\mathcal{K}$  iff both  $CL(\langle \mathcal{A} \cup \{loc(id, q, t)[0, \ell - \epsilon]\}, \mathcal{F})$  and  $CL(\langle \mathcal{A} \cup \{loc(id, q, t)[u + \epsilon, 1]\}, \mathcal{F})$  have no feasible solution, where  $\epsilon$  is as given in Theorem 10.*

Finally, we state the following corollary of Theorems 9 and 10 concerning cases where computing selection queries is tractable. Although the result is stated assuming that only binary std-formulas are in the PST KB, it holds for the more general case where both unary and binary std-formulas are in the PST KB since, using the approach suggested in the proof of Theorem 5, we can assume that unary std-formulas are encoded as PST atoms.

**Corollary 8.** *Let  $\mathcal{K} = \langle \mathcal{A}, \mathcal{F} \rangle$  be a PST KB where  $\mathcal{F}$  consists of binary std-formulas only. If  $\mathcal{G}$  is clique-acyclic or simple clique-cyclic, then deciding whether  $\langle id, t \rangle$  is an optimistic/cautious answer to selection query  $Q = (?id, q, ?t, [\ell, u])$  is in PTIME.*

## 5. Domain Enlargement

Up to this point we have assumed that the three domains  $ID = \{id_1, \dots, id_m\}$ ,  $T = [0, 1, \dots, tmax]$ , and  $Space = \{p_1, \dots, p_n\}$  are fixed. In this context we have investigated the consistency of a PST KB. In this section we investigate what happens to a PST KB if one or more domains are modified. The modification consists of possibly adding new time values, or spatial points, or objects. In fact, we are interested in what happens as we add arbitrarily many (but a finite number of) entities. In Section 5.1 we consider the case where we deal with a longer time period and add additional time values beyond  $tmax$ . In Section 5.2 we again deal with time but in this case we allow for a finer division of time values. Section 5.3 deals with the case where  $Space$  is enlarged, that is, new points are added. Finally, in Section 5.4 we deal both with additional objects as well as combinations of domain enlargement.

### 5.1 Extending Time Beyond $tmax$

We start with a PST KB  $\mathcal{K}$  using  $ID$ ,  $T$ , and  $Space$ . Suppose now that we extend  $T$  beyond  $tmax$ ; we write  $T' = [0, 1, \dots, tmax, \dots, t']$ . This means that the syntax must also be extended with the

constants  $tmax + 1, \dots, t'$ ; however,  $\mathcal{K}$  does not use any of the new constants. But the semantics is different now because each world  $w'$  must be a function  $w' : ID \times T' \rightarrow Space$ . We write  $W$  for the set of worlds using  $T$  and  $W'$  for the set of worlds using  $T'$ . The connection between  $W$  and  $W'$  is that every  $w \in W$  can be extended to many worlds in  $W'$  by choosing a point in  $Space$  for each new time value for each object. But from any one world  $w' \in W'$  there is a unique restriction to a world in  $W$ . Satisfaction for the worlds in  $W'$  is defined in the same way as satisfaction for the worlds in  $W$  (with  $W'$  substituted for  $W$ ).

Our interest in the semantics is due to the concept of interpretation where an interpretation assigns each world a probability. Let  $W = \{w_1, \dots, w_z\}$  where  $z = n^{m \times (tmax+1)}$  and let  $W' = \{w'_1, \dots, w'_y\}$  where  $y = n^{m \times (t'+1)}$ . For every  $w_i \in W$ , we define the set of worlds  $W'_i = \{w'_j | w'_j \in W' \text{ and } w'_j \text{ is an extension of } w_i\}$ . This means that the restriction of every  $w'_j \in W'_i$  is  $w_i$ . Next, let  $I_1$  be an interpretation on  $W$ . We call an interpretation  $I'_1$  an extension of  $I_1$  if for all  $i$ ,  $1 \leq i \leq z$ ,  $\sum_{w' \in W'_i} I'_1(w') = I_1(w_i)$ . That is, the probability value of  $I_1(w_i)$  is distributed among the extensions of  $w_i$  for each  $i$ . We can also go in the other direction. If  $I'_2$  is an interpretation on  $W'$ , we define its restriction to  $W$ ,  $I_2$ , by using the same formula, that is,  $I_2(w_i) = \sum_{w' \in W'_i} I'_2(w')$  for all  $i$ ,  $1 \leq i \leq z$ .

This framework allows us to investigate what happens to the consistency and inconsistency of  $\mathcal{K}$  as we go from  $T$  to  $T'$ . We start with inconsistency. We state the next theorem in two different forms because both are useful in various contexts.

**Theorem 12.** *Let  $\mathcal{K}$  contain only time values in  $T$ .*

- a) *If  $\mathcal{K}$  is inconsistent with time  $T$ , then  $\mathcal{K}$  remains inconsistent if  $T$  is expanded to  $T'$ .*
- b) *If  $T'$  is an expansion of  $T$  and  $\mathcal{K}$  is consistent with  $T'$  then it is consistent with  $T$ .*

*Proof.* As the second part is the contrapositive of the first, it suffices to prove just one part. We prove the second one. So assume that  $\mathcal{K}$ , which uses only time values in  $T$ , is consistent when  $T$  is expanded to  $T'$ . This means that there is an interpretation  $M'$  for  $\mathcal{K}$  (using  $W'$ ) that is a model of  $\mathcal{K}$ . We obtain an interpretation  $M$  (for  $W$ ) from  $M'$  using the formula given above, that is,  $M(w_i) = \sum_{w' \in W'_i} M'(w')$  for all  $i$ ,  $1 \leq i \leq z$ . We need to show that  $M$  is a model of  $\mathcal{K}$ .

Consider first the atomic formulas in  $\mathcal{A}$ . Clearly, for any world  $w_i$ ,  $w_i \models loc(id, r, t)$  iff  $w' \models loc(id, r, t)$  for all  $w' \in W'_i$ . Hence,  $\sum_{w | w \models loc(id, r, t)} M(w) = \sum_{w' | w' \models loc(id, r, t)} M'(w')$ . This takes care of  $\mathcal{A}$ .

Next, let  $f \in \mathcal{F}$ . Then  $\sum_{w' | w' \not\models f} M'(w') = 0$ . We need to show that also  $\sum_{w | w \not\models f} M(w) = 0$ . Let  $w_i \in W$  be a world such that  $w_i \not\models f$ . This means that there is a substitution  $\theta$  using only time values in  $T$  such that  $w_i \not\models \theta(f)$ . But  $\theta$  remains a substitution when  $T'$  is used for time. So for all  $w'_j \in W'_i$ ,  $w'_j \not\models \theta(f)$ . As  $M'$  is a model of  $\mathcal{K}$  and hence  $\mathcal{F}$ ,  $\sum_{w' | w' \not\models f} M'(w') = 0$  and the result follows from the definition of  $M$ .  $\square$

Next we consider the case where a PST KB defined using  $T$  is consistent. It turns out that consistency need not be preserved when  $T$  is expanded to  $T'$ . Consider the case of the following single integrity constraint:

$\forall X_1 X_2 Y_1 Z_1 \neg [loc(X_1, Y_1, Z_1) \wedge loc(X_2, Y_1, Z_1) \wedge X_1 \neq X_2 \wedge tmax < Z_1]$  and suppose that  $|ID| > |Space|$ . The std-formula states that two different objects cannot be in the same location after time  $tmax$ . The condition  $tmax < Z_1$  is false for all substitutions for  $Z_1$ ; hence the integrity constraint is automatically true. But once  $T$  is enlarged to any  $T' (> T)$ , say  $t' = tmax + 1$ , there is

an inconsistency because there are not enough points in *Space* for all the objects to occupy distinct points. We now show that for a given consistent  $\mathcal{K}$  we can always extend  $T = [0, \dots, tmax]$  to  $T^{big} = [0, \dots, tmax, \dots, tbig]$  such that if  $\mathcal{K}$  is consistent with  $T^{big}$  then  $\mathcal{K}$  remains consistent for any  $T'$  with  $T' = [0, \dots, tmax, \dots, t']$ . Essentially, we must make sure that for every time variable there is a substitution that makes each conjunct true in each  $f$ . For example, if the previous example is modified to

$\forall X_1 X_2 Y_1 Z_1 Z_2 \neg[loc(X_1, Y_1, Z_1) \wedge loc(X_2, Y_1, Z_2) \wedge X_1 \neq X_2 \wedge tmax < Z_1 \wedge Z_2 = Z_1 + 4]$  then we must have enough time values to include  $tmax + 5$ . We show how to do this systematically in the proof of the next theorem.

**Theorem 13.** *For a consistent PST KB  $\mathcal{K}$  defined with time values  $T = [0, \dots, tmax]$ , there is always a finite time value  $tbig \geq tmax$  that can be computed in linear time, such that if  $\mathcal{K}$  is consistent using  $T^{big} = [0, \dots, tmax, \dots, tbig]$  then  $\mathcal{K}$  is consistent for all  $T' = [0, \dots, tmax, \dots, t']$ .*

*Proof.* We obtain a time value  $t_f$  for each integrity constraint  $f \in \mathcal{F}$ . For each  $f$  we initialize  $t_f = tmax + 1$ . If  $f$  contains at most one time variable  $Z$  then we are done for  $f$ . So suppose that  $f$  has at least one conjunct with two distinct time variables. That must have one of the following four forms (where any  $+0$  is omitted): 1)  $Z_i + m = Z_j + n$ , 2)  $Z_i + m \neq Z_j + n$ , 3)  $Z_i + m < Z_j + n$ , or 4)  $Z_i + m \geq Z_j + n$ . The process involves adding a certain number to  $t_f$  for each such conjunct. For type 1) add  $|m - n|$ , for type 2) add 1, for type 3) add  $max(0, m - n + 1)$ , and for type 4) add  $max(0, n - m)$ . Then, by adding all the numbers for such conjuncts we obtain a  $t_f$  value for each  $f \in \mathcal{F}$ . Let  $tbig = max_{f \in \mathcal{F}}\{t_f\}$ . This process is linear in the size of  $\mathcal{F}$ .

We must show that if  $\mathcal{K}$  is consistent for  $T^{big}$  then it is consistent for any  $T' = [0, \dots, tmax, \dots, t']$ . If  $t' \leq tbig$  then consistency follows from Theorem 12b). So assume that  $tbig \leq t'$  and  $\mathcal{K}$  is consistent for  $T^{big}$ . This means that there is an interpretation  $M^{big}$  that is a model for  $\mathcal{K}$ . Just as in the proof of Theorem 12, the issue concerns only the constraints in  $\mathcal{F}$ . We write  $W^{big}$  for the set of all worlds using  $T^{big}$  and write  $w^b$  for an arbitrary world in  $W^{big}$ . For any  $f \in \mathcal{F}$ ,  $\sum_{w^b | w^b \not\models f} M^{big}(w^b) = 0$ . We obtain  $M'$  from  $M^{big}$  in an arbitrary way as long as the formula  $\sum_{w' \in W'_i} M'(w') = M^{big}(w_i^b)$  is satisfied. To show that  $M'$  is a model for  $\mathcal{K}$  with  $T'$ , we need to show that  $\sum_{w' | w' \not\models f} M'(w') = 0$ . So let  $w'_i \in W'$  be such that  $w'_i \not\models f$ . Then there is a substitution  $\theta'$  such that  $w'_i \not\models \theta'(f)$ . This  $\theta'$  may include any time values in  $T'$ . But our construction assures that there is already a substitution  $\theta^b$  using only values in  $T^{big}$  so that the world  $w_i^b$  which is the restriction of  $w'_i$  to  $T^{big}$  is such that  $w_i^b \not\models \theta^b(f)$ . Hence  $M'$  is a model for  $\mathcal{K}$  with  $T'$ .  $\square$

These results allow us to consider the case of arbitrarily large (finite) time.

**Definition 19.** *We call  $\mathcal{K}$  eventually consistent (resp. inconsistent) for time if there is an integer  $L$  such that  $\mathcal{K}$  is consistent (resp. inconsistent) for all  $T^\eta = [0, \dots, L, \dots, \eta]$ .*

**Corollary 9.** *Every  $\mathcal{K}$  is either eventually consistent for time or eventually inconsistent for time.*

*Proof.* If  $\mathcal{K}$  is inconsistent with time  $T$  then by Theorem 12a) it is eventually inconsistent for time. If  $\mathcal{K}$  is consistent with time  $T$  then there are two cases. If  $\mathcal{K}$  is still consistent with time  $T^{big}$  as computed in the proof of Theorem 13 then it is eventually consistent for time. Otherwise we conclude from Theorem 12a) that it is eventually inconsistent for time.  $\square$

Consider now how our results can be used for query answering when we expand time from  $T = [0, \dots, tmax]$  to  $T' = [0, \dots, tmax, \dots, t']$ . In order to avoid confusion we write  $\mathcal{K}$  where

the time is  $T$  and  $\mathcal{K}'$  when the time is  $T'$ . First, if  $\mathcal{K}$  is inconsistent, according to Theorem 12a),  $\mathcal{K}'$  is also inconsistent. So consider the case where  $\mathcal{K}$  is consistent (recall that queries are evaluated on consistent PST KBs). Note that when we check if the pair  $\langle id, t \rangle$  is an answer to  $Q$  with the expansion of  $T$  to  $T'$ , this considers also time values  $t$  where  $tmax < t \leq t'$ . Using Theorem 10 we obtain:

- $\langle id, t \rangle$  is an optimistic answer to  $Q$  w.r.t.  $\mathcal{K}'$  iff  $\langle \mathcal{A} \cup \{loc(id, q, t)[\ell, u]\}, \mathcal{F} \rangle$  is consistent using  $T'$ . We do not need to check separately that  $\mathcal{K}'$  is consistent because  $\mathcal{K}' \subset \langle \mathcal{A} \cup \{loc(id, q, t)[\ell, u]\}, \mathcal{F} \rangle$ .
- But to check for a cautious answer, we first check that  $\mathcal{K}'$  is consistent, and if so, then  $\langle id, t \rangle$  is a cautious answer to  $Q$  w.r.t.  $\mathcal{K}'$  iff  $\langle \mathcal{A} \cup \{loc(id, q, t)[0, \ell - \epsilon]\}, \mathcal{F} \rangle$  and  $\langle \mathcal{A} \cup \{loc(id, q, t)[u + \epsilon, 1]\}, \mathcal{F} \rangle$  are not consistent using  $T'$ , where  $\epsilon$  is as given in Theorem 10.

In a similar vein, Theorem 11 carries over also to the case of extended time.

## 5.2 Extending Time by More Frequent Time Values

In the second type of time extension we consider more frequent time values, so that there are more than  $tmax + 1$  time values between 0 and  $tmax$ . For illustration, we will use  $T' = [0, 0.5, 1, 1.5, \dots, tmax]$ . Again, the syntax must be changed as well to include the new constants 0.5, 1.5, and so on; however, the original  $\mathcal{K}$  is defined using  $T = [0, 1, \dots, tmax]$  and does not contain any new constant. The situation for worlds is similar to what happened in the previous subsection: every  $w : ID \times T \rightarrow Space$  is a unique restriction of many worlds,  $w'$ 's,  $w' : ID \times T' \rightarrow Space$ . Also, satisfaction for the worlds in  $W'$  using  $T'$  is defined the same way as for the worlds in  $W$  using  $T$ .

Our first result is the same for this type of time extension as for the previous one.

**Theorem 14.** *Let  $\mathcal{K}$  contain only time values in  $T$ .*

- a) *If  $\mathcal{K}$  is inconsistent with time  $T$ , then  $\mathcal{K}$  remains inconsistent if  $T$  is expanded to  $T'$ .*
- b) *If  $T'$  is an expansion of  $T$  and  $\mathcal{K}$  is consistent with  $T'$  then it is consistent with  $T$ .*

*Proof.* Analogous to the proof of Theorem 12. □

Next we consider the case where  $\mathcal{K}$  is consistent (using  $T$ ). Again, consistency need not be preserved in going from  $T$  to  $T'$ . Consider the case of the following std-formula:

$$\forall Y_1 Z_1 \neg [loc(id_1, Y_1, Z_1) \wedge Z_1 > 0 \wedge Z_1 < 1].$$

This formula states that  $id_1$  cannot be at any location between 0 and 1. The condition  $0 < Z_1 \wedge Z_1 < 1$  is false for all substitutions in  $T$ , hence the integrity constraint is true. But if we expand to  $T'$  then the substitution of 0.5 for  $Z_1$  makes the formula state that  $id_1$  cannot be at any location at time 0.5, making it inconsistent. However, we can extend  $T$  to more frequent time values  $T^{fr}$  such that if  $\mathcal{K}$  is consistent for  $T^{fr}$  then it will be consistent for any subdivision of the original time values.

**Theorem 15.** *For a consistent PST KB  $\mathcal{K}$  defined with time values  $T = [0, \dots, tmax]$ , there is always an integer value  $\chi$  that can be computed in linear time, such that if  $\mathcal{K}$  is consistent using  $T^{fr} = [0, \frac{1}{\chi}, \frac{2}{\chi}, \dots, 1, \frac{\chi+1}{\chi}, \dots, tmax]$  then  $\mathcal{K}$  remains consistent for every subdivision of  $T$ .*

*Proof.* The proof is analogous to the proof of Theorem 13. We calculate an integer  $\chi_f$  for each integrity constraint  $f \in \mathcal{F}$  and let  $\chi = \max_{f \in \mathcal{F}} \{\chi_f\}$ . The idea is to make sure that there are enough time values to allow any  $\gamma(\mathbf{Z})$  that can become true by some subdivision of time intervals to become true by substitution. In the worst case every time variable may require a new subdivision of the time values. Hence we can choose  $\chi_f = 2^k$  where  $k$  is the number of time variables in  $f$ . The rest of the proof is analogous to the proof of Theorem 13.  $\square$

As in the previous section where we wanted to consider arbitrarily large time values, now we consider dividing time intervals arbitrarily many times.

**Definition 20.** We call  $\mathcal{K}$  **divisionally consistent (resp. inconsistent) for time** if there is an integer  $L$  such that  $\mathcal{K}$  is consistent (resp. inconsistent) for all  $T_\chi = [0, \frac{1}{\chi}, \frac{2}{\chi}, \dots, tmax]$  where  $\chi > L$ .

**Corollary 10.** Every  $\mathcal{K}$  is either divisionally consistent for time or divisionally inconsistent for time.

*Proof.* Analogous to the proof of Corollary 9.  $\square$

Results analogous to the time expansion done in Section 5.1 hold here also for query answering and complexity.

### 5.3 Space Enlargement

In this subsection we consider what happens to the consistency or inconsistency of  $\mathcal{K}$  if *Space* is enlarged, say from  $Space = \{p_1, \dots, p_n\}$  to  $Space' = \{p_1, \dots, p_n, \dots, p_v\}$ . The change in semantics is different in this case than when we expanded  $T$ . If  $w : ID \times T \rightarrow Space$  is a world (for *Space*), it remains a world for *Space'*. So writing  $W$  for the set of worlds using *Space* and  $W'$  for the set of worlds using *Space'*, we obtain  $W \subset W'$ . There is no change in the definition of satisfaction, but the number of interpretations becomes greatly enlarged. Still, every interpretation  $I$  using *Space* can be extended to a unique  $I'$  using *Space'* by assigning  $I'(w') = 0$  for all  $w' \in W' \setminus W$ .

As in the case of time, we start with the case where  $\mathcal{K}$  is inconsistent. But unlike for time, when *Space* is enlarged  $\mathcal{K}$  may become consistent. Consider the simple example where  $\mathcal{K}$  consists of a single atom:  $a = loc(id_1, \{p_1, \dots, p_n\}, 1)[.2, .7]$ .  $\mathcal{K}$  is inconsistent. But if we add just a single point to *Space* so  $Space' = \{p_1, \dots, p_n, p_{n+1}\}$  then  $\mathcal{K}$  becomes consistent: for instance if  $w'_1$  is the world where  $w'_1(id_1, 1) = p_1$  and the other values are arbitrary, and  $w'_2$  is the world where  $w'_2(id_1, 1) = p_{n+1}$  and the other values are arbitrary, then assigning  $I'(w'_1) = 0.5$  and  $I'(w'_2) = 0.5$  with  $I'(w') = 0$  for all other worlds makes  $I'$  a model. A similar situation may occur with integrity constraints. Consider  $\mathcal{K}$  that contains a single std-formula:  $f : \forall X_1 Y_1 Z_1 \neg [loc(X_1, Y_1, Z_1) \wedge Y_1 \text{ ov } \{p_1, \dots, p_n\}]$  Since every region overlaps with *Space*,  $f$ , by itself, is inconsistent. But again, if we enlarge *Space* by one point  $p_{n+1}$  to *Space'* we can find a model  $I'$  as follows. Let  $I'(w') = 0$  for every world  $w'$  where  $w'(id, t) \in Space$  for some  $\langle id, t \rangle$ -pair, and  $I'(w_0) = 1$  where  $w_0(id, t) = p_{n+1}$  for all  $\langle id, t \rangle$  pairs. Then  $\sum_{w' | w' \neq f} I'(w) = 0$ ; hence the inconsistency is removed. However, suppose that in addition to  $f$ ,  $\mathcal{K}$  also contains the atom  $a = loc(id_1, \{p_1\}, 0)[1, 1]$  in  $\mathcal{A}$ , that is, now  $\mathcal{A} = \{a\}$  and  $\mathcal{F} = \{f\}$ . The PST atom  $a$  states that  $id_1$  must be at  $p_1$  at time 0. In that case we cannot make  $\mathcal{K}$  consistent by adding any number of points to *Space* because  $a$  and  $f$  are in conflict. Hence there is no general statement about what happens to an inconsistent PST KB when *Space* is extended to *Space'*.

Next we consider the case where  $\mathcal{K}$  is consistent (using  $Space$ ). In this case we can show that it remains consistent when  $Space$  is expanded to  $Space'$ .

**Theorem 16.** *If  $\mathcal{K}$  is consistent with  $Space$ , then  $\mathcal{K}$  remains consistent when  $Space$  is enlarged to  $Space'$ .*

*Proof.* Let  $M$  be a model for  $\mathcal{K}$  (using  $Space$ ).  $M$  assigns a probability  $M(w)$  to each world  $w$ . When  $Space$  is extended to  $Space'$ , many new worlds are added. We define  $M'(w) = M(w)$  if  $w$  is a world where only  $Space$  is used and  $M'(w) = 0$  otherwise. Thus by basically excluding the worlds using points other than  $Space$  in the sense that they are given probability 0,  $M'$  is a model of  $\mathcal{K}$  using  $Space'$ . Hence consistency is preserved.  $\square$

Again, just as in the case of time, we are interested in what happens for arbitrarily large finite spaces.

**Definition 21.** *We call  $\mathcal{K}$  eventually consistent (resp. inconsistent) for space if there is an integer  $L$  such that  $\mathcal{K}$  is consistent (resp. inconsistent) for all  $Space^\eta = \{p_1, \dots, p_n, \dots, p_L, \dots, p_\eta\}$ .*

**Corollary 11.** *Every  $\mathcal{K}$  is either eventually consistent for space or eventually inconsistent for space.*

*Proof.* There are two cases. We start with the first case where  $\mathcal{K}$  is consistent with  $Space$ . Then, by Theorem 16 it remains consistent for any larger  $Space'$ . Hence it is eventually consistent for space. Consider now the second case where  $\mathcal{K}$  is inconsistent with  $Space$ . There are two subcases. First, suppose that  $\mathcal{K}$  becomes consistent with some larger  $Space'$ . Then, just as in the first case, by Theorem 16  $\mathcal{K}$  is eventually consistent for space. The second subcase is where  $\mathcal{K}$  never becomes consistent no matter how  $Space$  is extended to  $Space'$ . This means that  $\mathcal{K}$  is eventually inconsistent for space.  $\square$

We showed earlier by examples that an inconsistent  $\mathcal{K}$  (with  $Space$ ) may become consistent if  $Space$  is enlarged. We can calculate a bound for the size of the needed space such that after that bound is reached the consistency or inconsistency does not change with further spatial enlargement.

**Theorem 17.** *For every inconsistent  $\mathcal{K}$  using  $Space = \{p_1, \dots, p_n\}$ , there is an explicit bound  $L$  that is tractable to compute such that if  $\mathcal{K}$  using  $Space' = \{p_1, \dots, p_L\}$  is inconsistent, it remains inconsistent for any enlargement of  $Space'$ .*

*Proof.* We already know from the above corollary that such an  $L$  exists. We now show how to compute it. First of all, if no enlargement of  $Space$  resolves the inconsistency of  $\mathcal{K}$  then  $\mathcal{K}$  is eventually inconsistent for space and we can choose  $L = n$ . Thus we need to deal in more detail only with the case where an enlargement of  $Space$  resolves the inconsistency of  $\mathcal{K}$ , that is,  $\mathcal{K}$  is eventually consistent for space. So we consider how adding points to  $Space$  can make  $\mathcal{K}$  consistent. There are three cases:

1. the inconsistency is due to  $\mathcal{A}$  alone,
2. the inconsistency is due to a combination of elements from both  $\mathcal{A}$  and  $\mathcal{F}$ , and
3. the inconsistency is due to  $\mathcal{F}$  only.

Consider first that in Case 1 an inconsistency must be due to atoms with some fixed  $id$  and  $t$  values. So we assume that we are dealing only with atoms using a specific pair of values:  $id$  and  $t$ . Adding a point or points to  $Space$  can resolve an inconsistency only if the atoms using  $id$  and  $t$  give  $Space$  a probability less than 1. For instance, let  $\mathcal{A} = \{loc(id, r_1, t, [0, 0.4]), loc(id, r_2, t, [0.1, 0.3])\}$  where  $r_1 \cup r_2 = Space$ . Then, enlarging  $Space$  to  $Space' = Space \cup \{p_{n+1}\}$  resolves the inconsistency. This is true in general, not just for this example, because once  $p_{n+1}$  is added to  $Space$ ,  $\{p_1, \dots, p_n\}$  becomes a proper subset of the relevant  $Space'$  and may consistently have probability less than 1. Now, consider that  $\mathcal{A}$  may have inconsistencies involving multiple pairs of  $id$  and  $t$  values. But that does not matter because the addition of the single point  $p_{n+1}$  to  $Space$  resolves all such inconsistencies. So in Case 1 we can choose  $L = n + 1$ .

In Case 2 an inconsistency of  $\mathcal{K}$  is due to a combination of elements from  $\mathcal{A}$  and  $\mathcal{F}$ . As every atom  $a$  in  $\mathcal{A}$  refers to a specific region  $r$  for a specific object  $id$  and time value  $t$ , for an inconsistency to occur with some  $f \in \mathcal{F}$ , some substitution of  $f$  must act as an atom (with probability interval  $[0, 0]$ ). For instance, let  $a = loc(id, r_1, t, [0, 0.4])$  and  $f = \forall X_1 \neg [loc(X_1, r_2, t)]$  where  $r_1 \cup r_2 = Space$ . The instance of  $f$  that causes the inconsistency is  $\neg [loc(id, r_2, t)]$  which has the same effect by the semantics as the atom  $loc(id, r_2, t, [0, 0])$ . This is true in general, not just for this example. In this way Case 2 reduces to Case 1 and we can choose  $L = n + 1$ .

The last case, Case 3, is where an inconsistency is due just to  $\mathcal{F}$ . The problem in this case is due to the fact that  $\mathcal{F}$  requires more points than are in  $Space$ . So we need to consider what can be expressed about the number of spatial points by std-formulas. This cannot be done by writing many spatial variables because if formula  $f$  has  $m$  spatial variables, say  $Y_1, \dots, Y_m$ , there is no way to express that they must all refer to different points. However, we can express that  $Space$  does not have enough points by writing  $f_0 = \forall X_1 Y_1 Z_1 \neg [loc(X_1, Y_1, Z_1) \wedge Y_1 \text{ov } Space]$ . Again, adding a single point to  $Space$  resolves the inconsistency. We can also write std-formulas that require at least a certain number of points. We can express the constraint that a point can be occupied by at most one object at one time by using the following 3 std-formulas:

$$f_1 = \forall X_1 Y_1 Z_1 Z_2 \neg [loc(X_1, Y_1, Z_1) \wedge loc(X_1, Y_1, Z_2) \wedge Z_1 \neq Z_2]$$

$$f_2 = \forall X_1 X_2 Y_1 Z_1 \neg [loc(X_1, Y_1, Z_1) \wedge loc(X_2, Y_1, Z_1) \wedge X_1 \neq X_2]$$

$$f_3 = \forall X_1 X_2 Y_1 Z_1 Z_2 \neg [loc(X_1, Y_1, Z_1) \wedge loc(X_2, Y_1, Z_2) \wedge X_1 \neq X_2 \wedge Z_1 \neq Z_2]$$

$f_1$  states that no object can be at the same point at two different times;  $f_2$  states that at no time can two objects be at the same point;  $f_3$  states that two different objects cannot be at the same point at two different times. Hence using std-formulas we can reference the number of objects times the number of time points. Let  $L = n + |ID| \times (tmax + 1)$ . Thus  $Space' = \{p_1, \dots, p_L\}$  consists of a point for each object and time point in addition to the original  $n$  points in  $Space$ , so that we can place objects in  $Space'$  even when none of the points in  $Space$  can be used and each object needs a new point at each time point. Clearly,  $L > n + 1$ , so this  $L$  works for Cases 1 and 2 as well. So this is the value of  $L$  that we choose for the tractable explicit bound. Thus if  $\mathcal{K}$  using  $Space' = \{p_1, \dots, p_L\}$  is still inconsistent, no points added to  $Space'$  can make it consistent.  $\square$

#### 5.4 Extending the Number of Objects or Several Entities

The last case where the number of constants may be increased is for  $ID$ . Since a world  $w$  is a function  $w : ID \times T \rightarrow Space$ , the expansion of objects is similar to the expansion of time, not

to the expansion of space. So the extension and restriction of interpretations going from an  $ID$  to a larger  $ID'$  is analogous to the addition of time values to  $T$ . Without getting into the details which are analogous to the cases for time given earlier, we can state that if  $\mathcal{K}$  is inconsistent with  $ID$ , it remains inconsistent if  $ID$  is expanded to  $ID'$ . To show that consistency need not be preserved, consider  $f = \forall X_1 X_2 \dots X_{m+1} \neg [\text{loc}(X_1, Y_1, Z_1) \wedge \text{loc}(X_2, Y_2, Z_2) \wedge \dots \wedge \text{loc}(X_{m+1}, Y_{m+1}, Z_{m+1}) \wedge X_1 \neq X_2 \wedge X_1 \neq X_3 \wedge \dots \wedge X_i \neq X_j \wedge \dots \wedge X_m \neq X_{m+1}]$  for  $i \neq j$ . Recalling that  $ID = \{id_1, \dots, id_m\}$  we find that  $\mathcal{K} = \{f\}$  is consistent with  $ID$  because it states that there cannot be more than  $m$  distinct objects. But it becomes inconsistent if  $ID$  is enlarged to  $ID' = \{id_1, \dots, id_m, id_{m+1}\}$ . But we can always find a  $q \geq m$  in linear time such that if  $\mathcal{K}$  is consistent for  $ID' = \{id_1, \dots, id_m, \dots, id_q\}$  then  $\mathcal{K}$  is consistent for any enlargement of  $ID'$ . Obtaining the value  $q$  is easier than obtaining  $tbig$  for time. It is just a matter of counting up the number of  $X$  variables in each  $f \in \mathcal{F}$  and taking the maximum such value, if bigger than  $m$ . Hence if we define eventually consistent (resp. inconsistent) for objects in the same way as for time, we obtain the result that every  $\mathcal{K}$  is either eventually consistent for objects or eventually inconsistent for objects.

So far we have only considered individual enlargements of either time values or points in space or objects. But we may be interested also in combining several types of extensions. As the expansion of time and objects are similar, we start by combining them.

**Definition 22.** We call  $\mathcal{K}$  using  $ID$ ,  $T$ , and  $Space$  **eventually consistent (resp. inconsistent) for objects and general time** if there are integers  $L_1$ ,  $L_2$ , and  $L_3$  such that  $\mathcal{K}$  is consistent (resp. inconsistent) for all  $ID_\iota = \{id_1, \dots, id_m, \dots, id_{L_1}, \dots, id_\iota\}$ ,  $Space = \{p_1, \dots, p_n\}$ , and  $T_\chi^\eta = [0, \frac{1}{\chi}, \frac{2}{\chi}, \dots, 1, \frac{\chi+1}{\chi}, \dots, tmax, \frac{\chi \cdot tmax+1}{\chi}, \dots, \dots, \eta]$  where  $\chi > L_2$  and  $\eta > L_3$ .

The combination of the expansion of objects and time both in magnitude and divisionally works essentially in the same way as the expansion of only one of these items. That is, if the set of all worlds for the original  $\mathcal{K}$  is  $W$  and the set for all worlds with the expansion to  $ID_\iota$  and  $T_\chi^\eta$  is  $W'$ , then for every world in  $W$  there is a set of extensions in  $W'$  and for every world in  $W'$  there is a unique world in  $W$  that is its restriction. The key issue is that  $Space$  remains unchanged.

Now let  $\mathcal{K}$  contain time values in  $T$ , regions in  $Space$  and objects in  $ID$ . Suppose that  $T$  is expanded to  $T'$  where the expansion may be both in magnitude and division and  $ID$  is expanded to  $ID'$  and find that  $\mathcal{K}$ , in the context of  $T'$  and  $ID'$  with  $Space$  unchanged, is consistent. Just as in the proof of Theorem 12 we start with a model  $M'$  for  $\mathcal{K}$  using  $T'$  and  $ID'$  and obtain the corresponding model  $M$  for  $\mathcal{K}$  using the starting  $T$  and  $ID$ . Thus, by the contrapositive, if  $\mathcal{K}$  is inconsistent with  $T$ ,  $Space$  and  $ID$ , it remains inconsistent with  $T'$ ,  $Space$  and  $ID'$ . This shows that any inconsistent  $\mathcal{K}$  is eventually inconsistent for objects and general time. In case  $\mathcal{K}$  is consistent with  $T$ ,  $Space$  and  $ID$  and becomes inconsistent with the expansion to  $T'$ ,  $Space$  and  $ID'$ , then again it is eventually inconsistent for objects and general time. The only alternative is for  $\mathcal{K}$  to remain consistent no matter how  $T$  and  $ID$  are expanded. Hence our earlier results can be put together as the following result.

**Theorem 18.** Every PST KB  $\mathcal{K}$  is either eventually consistent for objects and general time or it is eventually inconsistent for objects and general time.

However, the situation is different in the cases where we combine space with time or with objects as we show using an example.<sup>12</sup>

12. This example was suggested by one of the reviewers.

**Example 25.** Consider the PST KB  $\mathcal{K} = \langle \emptyset, \{f\} \rangle$  defined using  $ID = \{id_1, \dots, id_m\}$ ,  $T = [0, 1, \dots, tmax]$ , and  $Space = \{p_1, \dots, p_n\}$  with  $m \leq n$  where  $f$  is the std-formula  $f_2$  used in the proof of Theorem 17 which states that two different objects cannot be at the same point at the same time. Because of the condition  $m \leq n$ ,  $\mathcal{K}$  is consistent. Expand  $ID$  to  $ID' = \{id_1, \dots, id_m, \dots, id_{n+1}\}$  leaving  $T$  and  $Space$  unchanged and call the new KB  $\mathcal{K}'$ . Now there are too many objects for the number of points; hence  $\mathcal{K}'$  is inconsistent. Next, expand  $Space$  to  $Space'' = \{p_1, \dots, p_n, p_{n+1}\}$  while leaving  $ID'$  and  $T$  unchanged and call the new KB  $\mathcal{K}''$ .  $\mathcal{K}''$  is consistent.  $\square$

This example shows that we cannot claim eventual (in)consistency if both the number of objects and the number of points may increase as we can just continue the process indefinitely. An analogous result holds in case both the number of time values and the number of points may increase.

## 6. Related Work

We first discuss related works on classical probabilistic logic that does not have explicit spatial and temporal components. Then we discuss the relationship between our work and other spatio-temporal approaches. Finally, we relate our framework to object tracking.

### 6.1 Probabilistic Logic

As discussed in Section 3, a PST KB can be expressed in classical propositional logic (Hailperin, 1984; Nilsson, 1986; Paris, 1994), and in particular the consistency checking problem can be formulated in terms of Probabilistic Satisfiability (PSAT), whose first formulation is attributed to Boole (1854). After its presentation to the AI community by Nilsson (1986), the study of PSAT from the point of view of efficient algorithms and computational complexity was first addressed by Georgakopoulos et al. (1988), who showed that PSAT is in NP and NP-hard even for binary clauses. The tractable results identified by Georgakopoulos et al. concern the special case where each clause involves at most two literals (2PSAT) and the *graph of clauses* is *outerplanar*,<sup>13</sup> where the graph of clauses contains a vertex for each literal and two kinds of edges: i) an edge between each pair of literals built from the same propositional variable, and (ii) an edge for each pair of literals appearing in the same clause. We note that the PSAT formula  $\phi_{\mathcal{K}}$  encoding PST KB  $\mathcal{K}$  (see the proof of Theorem 1) contains more than two literals per clause even if we focus on binary std-formulas ( $\phi_{\mathcal{K}}$  becomes a 2PSAT formula only if we assume that  $Space$  consists of just two points).

The result provided by Georgakopoulos et al. relies on reducing 2PSAT to a tractable instance of 2MAXSAT (weighted maximum satisfiability problem with at most two literals per clause). Using this reduction and the result of Conforti and Cornuéjols (1992) on the tractability of problems that can be formulated as an integer program whose matrix is balanced, the following general result was provided by Conforti and Cornuéjols: PSAT is tractable for a *balanced set of clauses*, that is, for a set of clauses whose corresponding  $\{0, \pm 1\}$  clause-variable matrix is balanced. However, this result also doesn't help in finding tractable cases for PST KBs  $\mathcal{K}$  through reduction to PSAT  $\phi_{\mathcal{K}}$ , as considering three points in  $Space$  suffices to make the matrix corresponding to  $\phi_{\mathcal{K}}$  not balanced (it entails the presence of an odd cycle in the graph of clauses that, as observed by Andersen & Pre-

13. A graph is said to be outerplanar if it can be embedded in the plane so that all of its vertices lie on the same face. That is, it can be drawn in the plane without crossings in such a way that no vertex is totally surrounded by edges.

tolani, 2001, characterizes non-balanced matrices). Motivated by the fact that the tractable cases for PSAT identified by Georgakopoulos et al. and Conforti and Cornuéjols rely on using a polynomial time algorithm whose complexity is characterized by a high polynomial degree (specifically, the complexity is  $O(n^6 \log n)$  where  $n$  is the number of propositional variables), Andersen and Pretolani identified more efficient algorithms for the classes of balanced sets of clauses that can be represented by either hypertrees (where each hyperarc corresponds to the set of literals in a clause) or such that the *co-occurrence graph* is a partial k-tree (Bodlaender, 1998).

It is worth noting that none of the tractable cases identified in this paper can be derived from the results for PSAT described above by reducing  $\mathcal{K}$  to the PSAT formula  $\phi_{\mathcal{K}}$ . In fact, our tractable cases derive from the specific structure of PST KBs. On the other hand, our tractability results entail the tractability of PSAT instances that can be reduced to a tractable instance of the consistency checking problem. In particular, we can decide in polynomial time PSAT instances  $\phi_{\mathcal{K}}$  having the structure specified in our reduction from the consistency checking problem to PSAT, and such that the consistency checking problem for the corresponding PST KB  $\mathcal{K}$  turns out to be tractable.

In principle, the fact that there is a reduction from the consistency checking problem to PSAT enables well-known techniques for solving (general) instances of PSAT based on column generation (Kavvadias & Papadimitriou, 1990; Jaumard et al., 1991) to be used for addressing the consistency of PST KBs. The same holds for the problem of answering selection queries in PST KBs that, as shown in Section 4, can be addressed by solving suitable instances of the consistency checking problem. Recent approaches to solve PSAT using SAT (Finger & Bona, 2011) or Integer Linear Programming (Cozman & di Ianni, 2015) for column selection report experiments showing a phase transition behaviour (first observed by Finger and Bona for PSAT) depending on the fraction between the number of clauses and propositional variables as well as the number of probability assignments. Using these techniques PSAT instances with hundreds of propositional variables and clauses can be solved in reasonable time. However, we believe that reducing the consistency checking (or query answering) problem to PSAT and then applying these techniques as they are would not be a successful approach, as the number of propositional variables and clauses that would be generated would be huge even for small-size PST KBs. To avoid this problem, we conjecture that the specific structure of (PSAT formulas encoding) PST KBs could be exploited to devise more efficient techniques for solving the consistency checking and the query answering problems. In this regard, it would be interesting to investigate the connection between our framework and the emerging field of lifted probabilistic inference (Kersting, 2012), where the structure of FOL-constructs (such as indistinguishable individuals) is exploited to speed up the reasoning process, and see if some results can carry over into PST KBs.

The use of integrity constraints to encode domain knowledge has been studied by Lukasiewicz (1999, 2001) and Flesca, Furfaro, and Parisi (2014), for probabilistic frameworks which however do not explicitly deal with space and time. The problem of probabilistic deduction in the presence of conditional constraints over basic events was addressed by Lukasiewicz (1999), who identified tractable instances of probabilistic KBs, whose conditional constraints define conditional constraint trees, to support deduction over paths of premise/conclusions basic events. The problem of checking the consistency of relational probabilistic databases (where tuples can be viewed as basic events) in the presence of denial constraints was addressed by Flesca et al., who provided tractability results for constraints whose conflict hypergraph (Chomicki, Marcinkowski, & Staworko, 2004) is  $\gamma$ -acyclic (Fagin, 1983) as well as for a special kind of cyclic hypergraphs, which can encode neither clique-acyclic std-graphs nor simple clique-cyclic std-graphs. An important probabilistic logic pro-

gramming approach with conditional constraints was proposed by Lukasiewicz (2001), who studied the complexity of the satisfiability and the entailment problems for several types of formulas but without identifying tractable cases. Differently from the above-cited papers, the atomic information in our framework has a structure involving objects, space, and time, and thus atoms may also be intrinsically related because they are about the same object, or space, or time value.

## 6.2 Spatio-Temporal Approaches

Substantial work has been done on spatio-temporal logics (Gabelaia et al., 2005; Knapp et al., 2006) which combine spatial and temporal formalisms. This includes important contributions on qualitative spatio-temporal representation and reasoning (Muller, 1998; Wolter & Zakharyashev, 2000; Cohn & Hazarika, 2001), which focus on describing entities and qualitative relationships between them while dealing with discrete time. Cohn, Li, Liu, and Renz (2014) provided an up-to-date overview of the work done in the field of qualitative spatial reasoning, where recently the important problem of combining topological and directional information for extended spatial objects has been addressed. However, these works are not intended for reasoning about moving objects whose location at a given time is uncertain (they do not put probabilities into the mix). Yaman et al. (2004, 2005a, 2005b) focused on spatio-temporal logical theories that describe known plans of moving objects by sets of *go* atoms, each of them stating that an object can go from location  $L_1$  to  $L_2$ , leaving  $L_1$  and reaching  $L_2$  at some time in some intervals, and travelling with a speed in a given interval. Later, Parker et al. (2007b) extended this logic to include some probabilistic information about such plans. The SPOT framework of Parker et al. (2007a) further extended that work to uncertainty about where objects might be at a given time.

Past work on the SPOT framework investigated efficient algorithms for computing optimistic and cautious answers to selection queries (Parker et al., 2009; Parisi et al., 2010). The initial SPOT framework on which we build by adding integrity constraints has been implemented and tested on real US Navy databases containing ship location data (Parker et al., 2009; Parisi et al., 2010). Aggregate queries have been recently investigated by Grant et al. (2013), who proposed three semantics along with computational methods for evaluating them. As SPOT databases provide information on moving objects, one important aspect addressed by Parker et al. (2008) and then further investigated by Grant et al. (2010) is that of revising SPOT data so that information on these objects may be changed as objects move. Grant et al. (2010) proposed several strategies for revising SPOT data such as finding maximal consistent subsets, minimally modifying the spatial, temporal, object, or probability components of PST atoms. A full logic including negation, disjunction and quantifiers for managing SPOT data was recently proposed by Doder et al. (2013), who focused on finding sound and complete sets of axioms for several fragments of the logic. Grant, Parisi, and Subrahmanian (2013) provided a comprehensive survey on the SPOT framework where other related research is also reviewed.

While there is much work on spatio-temporal databases (Agarwal et al., 2003; Pelanis et al., 2006) and probabilistic spatio-temporal databases (Tao et al., 2005; Zhang, Chen, Jensen, Ooi, & Zhang, 2009; Zheng, Trajcevski, Zhou, & Scheuermann, 2011), these works mainly focus on devising indexing mechanisms and scaling query computation, instead of representing knowledge in a declarative fashion. In particular, Chung, Lee, and Chen (2009) use indexing to speed the computation of range queries and derive a PDF for the location of an object moving in a one-dimensional space by using its past moving behavior or the moving velocity distribution. Zhang et al. (2009)

provide a  $B^x$ -tree index which is a variant of a  $B^+$ -tree applicable to moving objects whose location and velocity are uncertain. Two types of pruning are introduced by Yang, Lu, and Jensen (2010) to efficiently solve queries asking for all sets of  $k$  objects that have at least a threshold probability of containing the  $k$  nearest objects to a given object. Dealing with a similar problem, Chen, Qin, and Liu (2010) propose a TPR-tree for indexing. Finally, Zheng et al. (2011) deal primarily with objects moving along road networks, and introduce an indexing mechanism for efficiently processing probabilistic range queries. However, none of these works systematically addresses the issue of considering integrity constraints over probabilistic spatio-temporal data.

### 6.3 Object Tracking

Object tracking is one of the most important problems in computer vision (Szeliski, 2010) where the consecutive positions of a tracked object are estimated as it moves in different frames of a video. Numerous approaches for object tracking have been proposed, mainly differing from each other on the type of object representation used (e.g., centroid, primitive geometric shapes), image features selected (e.g., colour, optical flow), and object detection method adopted (e.g., background subtraction, segmentation). However, the tracking algorithm to be chosen in a given application strongly depends on the application domain (Yilmaz, Javed, & Shah, 2006). Moreover, as object tracking algorithms may incur errors, due for instance to loss of information caused by projection of the 3D world on a 2D image, noise in images, partial and full object occlusions, the estimation of the position of tracked moving objects is inherently uncertain (even if a camera focuses on a fixed, specific area). Several important statistical methods for object tracking in computer vision (e.g., see Broida & Chellappa, 1986; Beymer & Konolige, 1999; Rosales & Sclaroff, 1999) are based on the well-known Kalman filter (1960) and its extensions to deal with the non-linear case, as well as on particle filtering (Kitagawa, 1987).

Filtering techniques have been extensively used for object tracking in the presence of sensors, not only cameras. As a matter of fact, object tracking has been extensively addressed in the more general setting where the position (i.e., the *state*) of one or more objects is estimated by a recursive Bayesian filter given the *measurements* up to that time coming from different kinds of sensors (including, for instance, radar, sonar, infrared, and other types of sensors possibly along with visual sensors) (Stone, Corwin, & Barlow, 1999). Basically, at each observed time point  $t$ , the output of the filter is a probability distribution (i.e., the *posterior*) on the position of the target object, which is computed by combining the motion updated at a time prior to  $t$  with the *likelihood* for the observation received at time  $t$ , where the likelihood represents the probability of each sensor measurement conditioned on the object position (Bar-Shalom, Kirubarajan, & Li, 2002).

The Kalman filter has been used for discrete-time estimation of continuous spatial positions of objects whose movement equations can be assumed to be linear with Gaussian noise. It has also been used successfully in non-linear systems by applying linearization and unscented transformation (Julier, Jeffrey, & Uhlmann, 2004). For discrete space and non-linear systems, particle filtering has been successfully used, providing a solution that can be applied to any state-space model and which generalizes the traditional Kalman filtering methods (Arulampalam, Maskell, Gordon, & Clapp, 2002). A general framework for particle filtering based on Sequential Importance Sampling with Resampling was proposed by Liu and Chen (1998), though a number of different types of particle filters exist and some have been shown to outperform others when used for particular applications (Arulampalam et al., 2002).

Differently from Kalman filtering, where the estimated position of an object at each observed time point is represented as a continuous distribution, particle filters are based on a histogram representation of the probability density, which is approximated by a finite number of *particles* (i.e., samples): each particle represents a position of space, and weights associated with particles (or the proportion of the number of particles) define a histogram of the probability distribution over space. This fits with the representation paradigm of PST KBs: as PST KBs allow us to represent, for each object and time point, a PDF over *Space* by defining a PST atom with a single valued probability interval (that is,  $\ell = u$ ) for each point in *Space*, it can be used easily to represent the output of object tracking techniques based on particle filtering. For filtering techniques returning a continuous distribution over *Space*, a discretization step should be applied.

While the output of object tracking techniques can be represented using PST KBs, there are important aspects that these techniques can deal with that PST KBs cannot do. In particular, filtering techniques use conditional independence to represent a PDF of an object's positions conditional on its positions at the previous time. PST KBs can encode the output of this inference process, but lack the expressive power to do that kind of inference. For instance, tracking techniques can represent the knowledge that if an object is in a region  $r_1$  at time  $t_1$  then it is probably in region  $r_2$  at time  $t_2$  with a probability depending on the time elapsed between  $t_1$  and  $t_2$ . Indeed trackers can rely on a motion model according to which the distribution for an object's location spreads out with the elapsed time since the last measurement: the distribution for the object's locations will be narrowly focused on locations near the measured position at  $t_1$  if  $t_2$  is close to  $t_1$ , but it will diffuse if  $t_2$  is faraway  $t_1$ . While tracking techniques can do this sort of things quite naturally, PST KBs can only capture some aspects of this behavior. For instance, we can express the fact that the object is probably in region  $r_1$  at  $t_1$ , and use an integrity constraint imposing that it can travel no more than  $d$  distance units in 1 time point. This would increase the probability of finding the object at time  $t_2$  in a region less than  $d$  units away from  $r_1$ , and it would decrease the probability of finding the object in a region farther away than  $d$  units. However, this is different from what can be inferred by tracking techniques using conditional independence.

On the other hand, tracking techniques do not combine very well with interval probabilities. In fact they typically return the PDF of an object's position at each observed time point. In contrast, using general PST atoms (with probability intervals), for each object and time point, all the (possibly infinite in number) PDFs that are compatible with the probability intervals specified by the atoms are succinctly represented. For instance, assuming  $Space = \{p_1, p_2\}$  and the PST atom  $loc(id, \{p_1\}, t)[0, 0.5]$ , all the PDFs  $f$  over *Space* assigning a probability  $f(p_1)$  in  $[0, 0.5]$  and  $f(p_2) = 1 - f(p_1)$  are represented. However, the set of PDFs represented can be restricted by using single valued probability intervals or by adding integrity constraints using std-formulas.

We note that the PST formalism allows us to impose integrity constraints over KBs that can be obtained by integrating position data coming from different sources. Consider for instance the integration of several PST KBs, each of them consisting of PST atoms encoding the output of an autonomous tracking system. The so-obtained integrated KB is still a PST KB, and integrity constraints can be used to express knowledge on the overall system, that could not be expressed considering the tracking systems separately. For instance, suppose we have an integrated PST KB consisting of the position data of monitored cars that are collected using black-box tracking systems installed over cars by insurance companies. Then std-formulas can be used to express correlations among monitored cars. For instance, knowing that in region  $r$  there is a licensed inspection station able to inspect at most  $k$  cars at the same time, we can impose the constraint that there cannot be

more than  $k$  cars in  $r$  at any time. This constraint would be meaningless if we considered the KB of each car tracking system separately, but it is useful for restricting the set of consistent interpretations of the PST KB obtained by integrating the several sources.

## 7. Summary

We believe that this is the first comprehensive paper that focuses systematically on knowledge representation in the form of integrity constraints for probabilistic spatio-temporal data.<sup>14</sup> The knowledge is represented both in the form of spatio-temporal atoms describing the location of objects in time with a probability interval as well as spatio-temporal denial formulas describing the integrity constraints the system must satisfy. Within this framework we investigated the consistency checking problem as well as the problem of answering selection queries in consistent PST KBs. Although both these problems turned out to be hard in the general case, we devised several sets of linear inequalities which allow us to decide consistency as well as answer queries by checking their feasibility. In addition, we identified different classes of spatio-temporal denial formulas for which both checking consistency and answering queries are tractable. Finally, we discussed an extension of the framework to arbitrarily large finite numbers of objects, time values, and points in space and showed that the behavior of consistency and inconsistency is uniform.

## 8. Future Work and Conclusion

There are further issues that can be investigated. Following Parisi and Grant (2014b), where we studied the problem of restoring the consistency of PST KBs of the form  $\mathcal{K}\langle\mathcal{A}, \emptyset\rangle$  (where the set of std-formulas is empty), we will consider the problem of repairing an inconsistent PST KB where inconsistency is due to the presence of std-formulas that are not satisfied. In this regard it would be interesting to devise methods for answering queries in inconsistent PST KBs. Recently there has been some research on probabilistic reasoning under inconsistency (Picado-Muiño, 2011; Thimm, 2013; Potyka & Thimm, 2014) that can help in this regard. It would also be interesting to look into the possibility of semantic query optimization for PST KBs, and to study the use of previous knowledge to efficiently check for consistency and process queries after updates.

Another direction for future work is the investigation of probabilistic std-formulas for expressing constraints that hold with a probability in a given interval. Intuitively, this would allow us to state for instance that two objects are not in the same region at the same time with probability greater than a given threshold, instead of stating that they cannot be in such a situation. This kind of probabilistic constraint could be expressed using pstd-formulas of the form  $f[\ell, 1]$  where  $f$  is an std-formula (i.e.,  $f[1, 1]$  captures the meaning of an std-formula). Then the change in semantics means changing the definition of model. Only the second part of Definition 6 has to be modified to  $\forall f \in \mathcal{F}, \sum_{w|w\models f} M(w) \in [\ell, 1]$ . Clearly, lower bounds on the complexity of consistency checking and query answering problems still hold for this extension. It is easy to check that the upper bound provided by Theorem 1 holds since a reduction to PSAT can still be provided by mapping pstd-formulas to clauses associated with a probability interval. As regards tractable cases, we conjecture that the results of Theorems 6 and 8 still hold for pstd-formulas of the form  $f[\ell, 1]$  if the right-hand side of inequality (4) of Definitions 15 and 17, respectively, is replaced with the lower probability bound  $\ell$  of the ground pstd-formula  $f[\ell, 1]$  generating the inequality. However, allowing general

14. This is a substantially revised and expanded version of the work by Parisi and Grant (2014a).

probability intervals associated with std-formulas introduces new issues in the semantics:  $f[\ell, u]$  would entail that  $\neg f[1 - u, 1 - \ell]$  holds. Providing a clear and intuitive semantics for this kind of std-formula, as well as for probabilistic std-formulas allowing probability intervals to be associated with each conjunct (instead of the whole formula), is deferred to future work.

The PST formalism is propositional even though the atoms have substantial content. We have added std-formulas for integrity constraints; these are a special class of first-order logic formulas. It would be interesting to consider other works on more general first-order probabilistic logics (Halpern, 1990; Lukasiewicz & Kern-Isberner, 1999; Kern-Isberner & Thimm, 2010). These logics were developed for a different purpose but an attempt could be made to represent spatio-temporal information in them. There are other works as well that could be enhanced with spatio-temporal information, such as Markov Logic (Richardson & Domingos, 2006), or Bayesian Logic Programs (Kersting & Raedt, 2007). In particular, Milch et al. (2005) introduces a first-order language called BLOG (Bayesian LOGic) for defining probability models over worlds with unknown objects and identity uncertainty, which finds a natural application in object tracking for unknown objects. It may be possible to find a generalization of the PST formalism that includes some of these concepts. In doing so, an aspect that we need to take care of is the fact that Markov logic and Bayesian logic programs deal with a unique probability distribution, while we deal with all the probability distributions compatible with the PST atoms and std-formulas.

Researchers in AI have been studying spatial and temporal reasoning for many years (Allen, 1984; Randell, Cui, & Cohn, 1992; Galton, 2009). An interesting project will be the incorporation of these concepts into the PST framework. The new syntax and semantics will include adding rules to the language. These additions will allow both for adding other types of information as well as new integrity constraints. For instance, using the concepts of qualitative direction and orientation proposed for spatial reasoning (Galton, 2009) would allow us to explicitly represent knowledge about the region toward which an object is moving. Another important concept needed in many applications is the explicit representation of qualitative and quantitative distance between objects as well as of information about the speed (i.e., maximum or average) of objects. Additional structured information about objects such as the type (e.g., vessel, vehicle, person, etc.) would be in general useful to exhaustively reason about moving objects. However, depending on the addition made to increase the expressive power of the extended framework, important consequences on the complexity of the consistency checking problem may arise. Spatial and temporal aspects of formalisms for qualitative spatial and temporal reasoning are more expressive than what we have in our framework (Gabelaia et al., 2005; Knapp et al., 2006). The trade-off between expressiveness and complexity within the hierarchy of formalisms obtained by combining well-known spatial and temporal logics is analyzed by Gabelaia et al. (2005), where it is shown that the complexity of the satisfiability problem for spatio-temporal logics (not dealing with probabilities) can vary from NP-complete to undecidable. Using these formalisms in the PST framework may drastically increase the computational complexity of the problems studied in this paper. Nevertheless, we believe that some attempt should later be made to include even simpler concepts of qualitative spatio-temporal reasoning to the PST framework, particularly trying to exploit restrictions recently studied by Huang, Li, and Renz (2013) to identify tractable fragments.

In this paper, we proposed a framework where four features of moving objects are taken into account: the spatial component, the temporal component, the inherent uncertainty of acquired data, and integrity constraints from the application domain. The expressiveness of these features could be improved to represent additional knowledge that may be of interest in practical applications, partic-

ularly using ideas from papers dealing with spatio-temporal reasoning. In the endeavor of a concrete implementation of our framework, ideas from the PostGIS spatial database system which extends PostgreSQL would be useful. We believe it will be worthwhile later to incorporate more concepts in the PST framework. We took a first step in extending the PST framework by the addition of integrity constraints, and hope that researchers will use our work as a starting point in investigations of the important role of integrity constraints in probabilistic spatio-temporal knowledge bases.

## Acknowledgments

We wish to thank the referees for numerous helpful comments that helped us in substantially improving this paper.

## References

- Agarwal, P. K., Arge, L., & Erickson, J. (2003). Indexing moving points. *J. Comput. Syst. Sci.*, 66(1), 207–243.
- Ahson, S. A., & Ilyas, M. (2010). *Location-Based Services Handbook: Applications, Technologies, and Security*. CRC Press, Hoboken, NJ.
- Akdere, M., Cetintemel, U., Riondato, M., Upfal, E., & Zdonik, S. B. (2011). The case for predictive database systems: Opportunities and challenges. In *Proceedings of the 5th Biennial Conference on Innovative Data Systems Research (CIDR)*, pp. 167–174.
- Allen, J. F. (1984). Towards a general theory of action and time. *Artif. Intell.*, 23(2), 123–154.
- Andersen, K. A., & Pretolani, D. (2001). Easy cases of probabilistic satisfiability. *Ann. Math. Artif. Intell.*, 33(1), 69–91.
- Arulampalam, M. S., Maskell, S., Gordon, N. J., & Clapp, T. (2002). A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2), 174–188.
- Bar-Shalom, Y., Kirubarajan, T., & Li, X.-R. (2002). *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc., New York, NY, USA.
- Bayir, M. A., Demirbas, M., & Eagle, N. (2010). Mobility profiler: A framework for discovering mobility profiles of cell phone users. *Pervasive and Mobile Computing*, 6(4), 435 – 454.
- Beymer, D., & Konolige, K. (1999). Real-time tracking of multiple people using continuous detection. In *Proceedings of the Workshop on Frame-rate Applications, Methods and Experiences with Regularly Available Technology and Equipment (FRAME-RATE)*, in conjunction with the 7th IEEE International Conference on Computer Vision (ICCV).
- Bodlaender, H. L. (1998). A partial  $k$ -arboretum of graphs with bounded treewidth. *Theor. Comput. Sci.*, 209(1-2), 1–45.
- Boole, G. (1854). *An Investigation of the Laws of Thought on Which are Founded the Mathematical Theories of Logic and Probabilities*. Macmillan, London.
- Broida, T. J., & Chellappa, R. (1986). Estimation of object motion parameters from noisy images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 8(1), 90–99.

- Chen, Y.-F., Qin, X.-L., & Liu, L. (2010). Uncertain distance-based range queries over uncertain moving objects. *J. Comput. Sci. Technol.*, 25(5), 982–998.
- Chomicki, J., Marcinkowski, J., & Staworko, S. (2004). Computing consistent query answers using conflict hypergraphs. In *Proceedings of the 2004 ACM CIKM International Conference on Information and Knowledge Management (CIKM)*, pp. 417–426.
- Chung, B. S. E., Lee, W.-C., & Chen, A. L. P. (2009). Processing probabilistic spatio-temporal range queries over moving objects with uncertainty. In *Proceedings of the 12th International Conference on Extending Database Technology (EDBT)*, pp. 60–71.
- Cohn, A. G., & Hazarika, S. M. (2001). Qualitative spatial representation and reasoning: An overview. *Fundam. Inform.*, 46(1-2), 1–29.
- Cohn, A. G., Li, S., Liu, W., & Renz, J. (2014). Reasoning about topological and cardinal direction relations between 2-dimensional spatial objects. *J. Artif. Intell. Res. (JAIR)*, 51, 493–532.
- Conforti, M., & Cornuéjols, G. (1992). A class of logic problems solvable by linear programming. In *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 670–675.
- Cozman, F. G., & di Ianni, L. F. (2015). Probabilistic satisfiability and coherence checking through integer programming. *Int. J. Approx. Reasoning*, 58, 57–70.
- Doder, D., Grant, J., & Ognjanović, Z. (2013). Probabilistic logics for objects located in space and time. *J. of Logic and Computation*, 23(3), 487–515.
- Fagin, R. (1983). Degrees of acyclicity for hypergraphs and relational database schemes. *Journal of the ACM*, 30(3).
- Finger, M., & Bona, G. D. (2011). Probabilistic satisfiability: Logic-based algorithms and phase transition. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 528–533.
- Flesca, S., Furfaro, F., & Parisi, F. (2014). Consistency checking and querying in probabilistic databases under integrity constraints. *J. Comput. Syst. Sci.*, 80(7), 1448–1489.
- Gabelaia, D., Kontchakov, R., Kurucz, Á., Wolter, F., & Zakharyashev, M. (2005). Combining spatial and temporal logics: Expressiveness vs. complexity. *J. Artif. Intell. Res.*, 23, 167–243.
- Galton, A. (2009). Spatial and temporal knowledge representation. *Earth Science Informatics*, 2(3), 169–187.
- Georgakopoulos, G. F., Kavvadias, D. J., & Papadimitriou, C. H. (1988). Probabilistic satisfiability. *J. Complexity*, 4(1), 1–11.
- Grant, J., Molinaro, C., & Parisi, F. (2013). Aggregate count queries in probabilistic spatio-temporal databases. In *Proceedings of the 7th International Conference on Scalable Uncertainty Management (SUM)*, pp. 255–268.
- Grant, J., Parisi, F., Parker, A., & Subrahmanian, V. S. (2010). An agm-style belief revision mechanism for probabilistic spatio-temporal logics. *Artif. Intell.*, 174(1), 72–104.
- Grant, J., Parisi, F., & Subrahmanian, V. S. (2013). Research in probabilistic spatiotemporal databases: The SPOT framework. In *Advances in Probabilistic Databases for Uncertain Information Management*, Vol. 304 of *Studies in Fuzziness and Soft Computing*, pp. 1–22. Springer.

- Hailperin, T. (1984). Probability logic. *Notre Dame Journal of Formal Logic*, 25(3), 198–212.
- Halpern, J. Y. (1990). An analysis of first-order logics of probability. *Artif. Intell.*, 46(3), 311–350.
- Hammel, T., Rogers, T. J., & Yetso, B. (2003). Fusing live sensor data into situational multimedia views. In *Proceedings of the 9th International Workshop on Multimedia Information Systems (MIS)*, pp. 145–156.
- Huang, J., Li, J. J., & Renz, J. (2013). Decomposition and tractability in qualitative spatial and temporal reasoning. *Artif. Intell.*, 195, 140–164.
- Jaumard, B., Hansen, P., & de Aragão, M. P. (1991). Column generation methods for probabilistic logic. *ORSA Journal on Computing*, 3(2), 135–148.
- Julier, S. J., Jeffrey, & Uhlmann, K. (2004). Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92, 401–422.
- Jünger, M., Liebling, T., Naddef, D., Nemhauser, G., Pulleyblank, W., Reinelt, G., Rinaldi, G., & Wolsey, L. (Eds.). (2010). *50 Years of Integer Programming 1958-2008: From the Early Years to the State-of-the-Art*. Springer, Heidelberg.
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering*, 82(Series D), 35–45.
- Karbassi, A., & Barth, M. (2003). Vehicle route prediction and time of arrival estimation techniques for improved transportation system management. In *Proceedings of the 2013 IEEE Intelligent Vehicles Symposium*, pp. 511–516.
- Karimi, H. A. (2013). *Advanced location-based technologies and services*. CRC Press, Hoboken, NJ.
- Kavvadias, D. J., & Papadimitriou, C. H. (1990). A linear programming approach to reasoning about probabilities. *Ann. Math. Artif. Intell.*, 1, 189–205.
- Kern-Isberner, G., & Thimm, M. (2010). Novel semantical approaches to relational probabilistic conditionals. In *Proceedings of the 12th International Conference on Principles of Knowledge Representation and Reasoning (KR)*.
- Kersting, K., & Raedt, L. D. (2007). Bayesian logic programming: Theory and tool. In Getoor, L., & Taskar, B. (Eds.), *An Introduction to Statistical Relational Learning*. MIT Press.
- Kersting, K. (2012). Lifted probabilistic inference. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI)*, pp. 33–38.
- Kitagawa, G. (1987). Non-gaussian state-space modeling of nonstationary time series. *Journal of the American Statistical Association*, 82(400), 1032–1041.
- Knapp, A., Merz, S., Wirsing, M., & Zappe, J. (2006). Specification and refinement of mobile systems in msla and mobile uml. *Theor. Comput. Sci.*, 351(2), 184–202.
- Kurkovsky, S., & Harihar, K. (2006). Using ubiquitous computing in interactive mobile marketing. *Personal Ubiquitous Comput.*, 10(4), 227–240.
- Li, S. Z., & Jain, A. K. (Eds.). (2011). *Handbook of Face Recognition, 2nd Edition*. Springer.
- Liu, J. S., & Chen, R. (1998). Sequential monte carlo methods for dynamic systems. *Journal of the American Statistical Association*, 93, 1032–1044.

- Lukasiewicz, T. (2001). Probabilistic logic programming with conditional constraints. *ACM Trans. on Computational Logic*, 2(3), 289–339.
- Lukasiewicz, T. (1999). Probabilistic deduction with conditional constraints over basic events. *J. Artif. Intell. Res. (JAIR)*, 10, 199–241.
- Lukasiewicz, T., & Kern-Isberner, G. (1999). Probabilistic logic programming under maximum entropy. In *Proceedings of the 5th European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU)*, pp. 279–292.
- MarketsandMarkets (2014). Location Based Services (LBS) Market (Mapping, Discovery and Infotainment, Location Analytics, Leisure and Social Networking, Location Based Advertising, Augmented Reality and Gaming, Tracking) - Worldwide Forecasts and Analysis (2014 - 2019). <http://www.marketsandmarkets.com/Market-Reports/location-based-service-market-96994431.html>.
- Milch, B., Marthi, B., Russell, S. J., Sontag, D., Ong, D. L., & Kolobov, A. (2005). BLOG: probabilistic models with unknown objects. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1352–1359.
- Mittu, R., & Ross, R. (2003). Building upon the coalitions agent experiment (coax) - integration of multimedia information in gccs-m using impact. In *Proceedings of the 9th International Workshop on Multimedia Information Systems (MIS)*, pp. 35–44.
- Muller, P. (1998). A qualitative theory of motion based on spatio-temporal primitives. In *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pp. 131–143.
- Nilsson, N. J. (1986). Probabilistic logic. *Artif. Intell.*, 28(1), 71–87.
- Papadimitriou, C. H., & Steiglitz, K. (1982). *Combinatorial optimization: algorithms and complexity*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- Papadimitriou, C. M. (1994). *Computational complexity*. Addison-Wesley, Reading, Massachusetts.
- Paris, J. (1994). *The Uncertain Reasoner's Companion: A Mathematical Perspective*. Cambridge University Press.
- Parisi, F., & Grant, J. (2014a). Integrity constraints for probabilistic spatio-temporal knowledgebases. In *Proceedings of the 8th International Conference on Scalable Uncertainty Management (SUM)*, pp. 251–264.
- Parisi, F., & Grant, J. (2014b). Repairs and consistent answers for inconsistent probabilistic spatio-temporal databases. In *Proceedings of the 8th International Conference on Scalable Uncertainty Management (SUM)*, pp. 265–279.
- Parisi, F., Parker, A., Grant, J., & Subrahmanian, V. S. (2010). Scaling cautious selection in spatial probabilistic temporal databases. In *Methods for Handling Imperfect Spatial Information*, Vol. 256 of *Studies in Fuzziness and Soft Computing*, pp. 307–340. Springer.
- Parisi, F., Sliva, A., & Subrahmanian, V. S. (2013). A temporal database forecasting algebra. *Int. J. of Approximate Reasoning*, 54(7), 827–860.
- Parker, A., Infantes, G., Grant, J., & Subrahmanian, V. S. (2009). SPOT databases: Efficient consistency checking and optimistic selection in probabilistic spatial databases. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 21(1), 92–107.

- Parker, A., Infantes, G., Subrahmanian, V. S., & Grant, J. (2008). An AGM-based belief revision mechanism for probabilistic spatio-temporal logics. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI)*, pp. 511–516.
- Parker, A., Subrahmanian, V. S., & Grant, J. (2007a). A logical formulation of probabilistic spatial databases. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 19(11), 1541–1556.
- Parker, A., Yaman, F., Nau, D. S., & Subrahmanian, V. S. (2007b). Probabilistic go theories. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 501–506.
- Pelani, M., Saltenis, S., & Jensen, C. S. (2006). Indexing the past, present, and anticipated future positions of moving objects. *ACM Trans. Database Syst.*, 31(1), 255–298.
- Petrova, K., & Wang, B. (2011). Location-based services deployment and demand: a roadmap model. *Electronic Commerce Research*, 11(1), 5–29.
- Picado-Muiño, D. (2011). Measuring and repairing inconsistency in probabilistic knowledge bases. *Int. J. Approx. Reasoning*, 52(6), 828–840.
- Potyka, N., & Thimm, M. (2014). Consolidation of probabilistic knowledge bases by inconsistency minimization. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*, pp. 729–734.
- Randell, D. A., Cui, Z., & Cohn, A. G. (1992). A spatial logic based on regions and connection. In *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pp. 165–176.
- Richardson, M., & Domingos, P. (2006). Markov logic networks. *Mach. Learn.*, 62(1-2), 107–136.
- Rosales, R., & Sclaroff, S. (1999). 3D trajectory recovery for tracking multiple objects and trajectory guided recognition of actions. In *Proceedings of the 6th Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2117–2123.
- Southey, F., Loh, W., & Wilkinson, D. F. (2007). Inferring complex agent motions from partial trajectory observations. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2631–2637.
- Stone, L. D., Corwin, T. L., & Barlow, C. A. (1999). *Bayesian Multiple Target Tracking* (1st edition). Artech House, Inc., Norwood, MA, USA.
- Szeliski, R. (2010). *Computer Vision: Algorithms and Applications*. Springer-Verlag New York, Inc., New York, NY, USA.
- Tao, Y., Cheng, R., Xiao, X., Ngai, W. K., Kao, B., & Prabhakar, S. (2005). Indexing multi-dimensional uncertain data with arbitrary probability density functions. In *Proceedings of the 31st International Conference on Very Large Data Bases (VLDB)*, pp. 922–933.
- Thimm, M. (2013). Inconsistency measures for probabilistic logics. *Artif. Intell.*, 197, 1–24.
- Wolter, F., & Zakharyashev, M. (2000). Spatio-temporal representation and reasoning based on rcc-8. In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pp. 3–14.

- Yaman, F., Nau, D. S., & Subrahmanian, V. S. (2004). A logic of motion. In *Proceedings of the 9th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pp. 85–94.
- Yaman, F., Nau, D. S., & Subrahmanian, V. S. (2005a). Going far, logically. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 615–620.
- Yaman, F., Nau, D. S., & Subrahmanian, V. (2005b). A motion closed world assumption. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 621–626.
- Yang, B., Lu, H., & Jensen, C. S. (2010). Probabilistic threshold k nearest neighbor queries over moving objects in symbolic indoor space. In *Proceedings of the 13th International Conference on Extending Database Technology (EDBT)*, pp. 335–346.
- Yilmaz, A., Javed, O., & Shah, M. (2006). Object tracking: A survey. *ACM Comput. Surv.*, 38(4).
- Zhang, M., Chen, S., Jensen, C. S., Ooi, B. C., & Zhang, Z. (2009). Effectively indexing uncertain moving objects for predictive queries. *Proceedings of the VLDB Endowment (PVLDB)*, 2(1), 1198–1209.
- Zheng, K., Trajcevski, G., Zhou, X., & Scheuermann, P. (2011). Probabilistic range queries for uncertain trajectories on road networks. In *Proceedings of the 14th International Conference on Extending Database Technology (EDBT)*, pp. 283–294.