
A Rewriting Framework and Logic for Activities Subject to Regulations

Max Kanovich, *Queen Mary, University of London, UK.*
E-mail: mik@eecs.qmul.ac.uk

Tajana Ban Kirigin, *University of Rijeka, Croatia.*
E-mail: bank@math.uniri.hr

Vivek Nigam, *Federal University of Paraíba, Brazil.*
E-mail: vivek.nigam@gmail.com

Andre Scedrov, *University of Pennsylvania, USA.*
E-mail: scedrov@math.upenn.edu

Carolyn Talcott, *SRI International, USA. E-mail: clt@csl.sri.com*

Ranko Perovic, *Senior Clinical Trial Specialist, USA.*
Email: perovicrankomd@gmail.com

Abstract

Activities such as clinical investigations or financial processes are subject to regulations to ensure quality of results and avoid negative consequences. Regulations may be imposed by multiple governmental agencies as well as by institutional policies and protocols. Due to the complexity of both regulations and activities there is great potential for violation due to human error, misunderstanding, or even intent. Executable formal models of regulations, protocols, and activities can form the foundation for automated assistants to aid planning, monitoring, and compliance checking. We propose a model based on multiset rewriting where time is discrete and is specified by timestamps attached to facts. Actions, as well as initial, goal and critical states may be constrained by means of relative time constraints. Moreover, actions may have non-deterministic effects, *i.e.*, they may have different outcomes whenever applied. We present a formal semantics of our model based on focused proofs of linear logic with definitions. Furthermore, we demonstrate how specifications in our model can be straightforwardly mapped to the rewriting logic language Maude, and how one can use existing techniques to improve performance. We also determine the computational complexity of various planning problems. Plan compliance problem, for example, is the problem of finding a plan that leads from an initial state to a desired goal state without reaching any undesired critical state. We consider all actions to be balanced, *i.e.*, their pre and post-conditions have the same number of facts. Under this assumption on actions, we show that the plan compliance problem is PSPACE-complete when all actions have only deterministic effects and is EXPTIME-complete when actions may have non-deterministic effects. Finally, we show that the restrictions on the form of actions and time constraints taken in the specification of our model are necessary for decidability of the planning problems.

1 Introduction

- 2 Regulations are commonly used to set the rules of conduct of numerous activities in order
- 3 to ensure quality of results and avoid negative consequences. For example, while carrying
- 4 out a clinical investigation (CI)—that is, a set of procedures in medical research and drug

2 A Rewriting Framework and Logic for Activities Subject to Regulations

5 development, to test a new drug or other intervention on human subjects, it is important
6 that conclusive data is collected and that the health of the subjects participating in the CI is
7 not compromised. In order to collect the most conclusive data, for instance, drug samples
8 have to be taken and all the necessary tests have to be carried out in well defined periods
9 of time. Moreover, since these experiments might compromise the health of subjects, CIs
10 are rigorously regulated by policies elaborated by governmental agencies such as the Food
11 and Drug Administration (FDA) [16]. These regulations require prompt action whenever a
12 serious and unexpected problem with any subject is reported. In the current state of affairs,
13 there is little to almost no automation in the management of CIs and therefore the process is
14 prone to human error. As described in [32], there is plenty of room for the use of automated
15 assistants to help reduce human mistakes from happening. For instance, a computer assistant
16 can automatically generate plans that guide the clinical staff on how a CI has to be carried out.
17 An assistant can also monitor the execution of a CI and signal alarms whenever a deviation to
18 the specification is detected.

19 This paper proposes a rewriting framework that can be used to specify collaborative systems,
20 such as CIs, and can be used as the foundation for building automated assistants. Our model is
21 an extension of the systems used for modeling collaborative systems proposed in [25] with
22 explicit time. An important feature of our model is that its specifications can be directly written
23 and executed in Maude [9], a powerful tool based on rewrite logic [29]. For more details see
24 [32] where we address implementation.

25 A second feature of our framework is that its specifications can mention time explicitly.
26 Time is often a key component used in policies specifying the rules and the requirements of a
27 collaboration. For a correct collaboration and to achieve a common goal, participants should
28 usually follow strict deadlines and should have quick reactions to some (unexpected) events.
29 For instance, the paragraph 312.32 on Investigational New Drug Application (IND) safety
30 [16] includes explicit time intervals that must be followed in case of any unexpected, serious
31 or life-threatening adverse drug experience: (The emphasis in the text below is ours.)

32 “(c) IND safety reports

33 (1) Written reports –(i) The sponsor shall notify FDA and all participating investigators
34 in a written IND safety report of: (A) Any adverse experience associated with the use
35 of the drug that is both serious and unexpected; [· · ·] Each notification shall be made
36 as soon as possible and *in no event later than 15 calendar days* after the sponsor’s
37 initial receipt of the information [· · ·]

38 (2) Telephone and facsimile transmission safety reports. The sponsor shall also no-
39 tify FDA by telephone or by facsimile transmission of any unexpected fatal or life-
40 threatening experience associated with the use of the drug as soon as possible but *in no*
41 *event later than 7 calendar days* after the sponsor’s initial receipt of the information.”

42 The above clause explicitly mentions two different time intervals. The first one specifies that a
43 detailed safety report must be sent to the FDA within 15 days after a serious and unexpected
44 event is detected, while the second specifies the obligation of notifying FDA of such an event
45 within 7 days.

46 In order to accommodate explicit time, we attach to facts a *natural number* called *timestamp*.
47 Timestamps can be used in different ways depending on the system being modeled. In the
48 example above, the timestamp t of the fact $Dose(id)@t$ could denote that the subject with
49 anonymous identification number id received a dose at time t . Alternatively, the timestamp, t_2 ,
50 of the fact $Deadline@t_2$ could denote the time of when some activity should end. Moreover,

51 we keep track of time by assuming a discrete global time, using the special fact $Time@t$
 52 that denotes that the current time is t . The global time advances by replacing $Time@t$ by
 53 $Time@(t + 1)$.

54 Agents change the state of the system by performing actions. In order to specify the type of
 55 time requirements illustrated above, a set of *time constraints* may be attached to actions. This
 56 set acts as a guard of the action, *i.e.* the action can only be applied if its time constraints are
 57 satisfied. Formally, a time constraint is a comparison involving exactly two timestamps, *e.g.*,
 58 $T_1 \leq T_2 + 7$ (see Eq. 2.1).

59 Besides allowing guards with time constraints, we also allow actions to have non-deter-
 60 ministic effects. In particular, actions are allowed to have a finite number of post-conditions
 61 specifying a finite number of possible resulting states. These actions are useful when specifying
 62 systems, such as CIs, containing actions that may lead to different outcomes, but it is not
 63 certain beforehand which one of the outcomes will actually occur. For instance, when carrying
 64 out a blood test for the presence of some substance, it is not a priori clear what the test result
 65 will be. Nevertheless, one can classify any result as either *positive* or *negative*. Depending on
 66 this result, one would need to take a different set of future actions. For example, if the blood
 67 test is positive, then one might not be suitable for participating as a subject in a particular CI,
 68 but may be suitable for other CIs. We classify actions that have more than one outcome as
 69 *branching actions*.

70 Finally, in collaborative systems agents collaborate in order to achieve a common goal, but
 71 they should also avoid *critical states* that, for example, violate policies. An example of a goal
 72 state for CIs would be to collect conclusive data without compromising the health of subjects,
 73 while a critical state would be a state that violates the FDA policies. In our model, critical,
 74 goal and initial states can also mention time explicitly by using time constraints.

75 This paper's contributions are the following:

- 76 1. Timed local state transition systems are specified in order to formalize systems with
 77 explicit time. This specification takes necessary restrictions on the type of actions and time
 78 constraints so that explicit time requirements are expressible in the system, but at the same
 79 time it is precise with respect to complexity results of the associated planning problems so
 80 that we provide decidability.
- 81 2. We determine the complexity of the *plan compliance* problem [25], that is, the problem of
 82 determining whether there is a plan where the collaboration achieves the common goal
 83 and in the process no critical state is reached. It has been shown that the plan compliance
 84 problem is undecidable in general [21]. However, we get decidability in the important
 85 case when all actions are *balanced*, *i.e.*, pre and post-conditions of actions have the same
 86 number of facts. Intuitively, this restriction bounds the memory of agents, as they can
 87 remember at any point only a bounded number of facts. Additionally, we assume that
 88 the facts created by an action, that is, the new facts that appear in its post-condition,
 89 can only have timestamps of the form $T + d$, where T is the current global time and d
 90 a natural number. Under these two assumptions on actions, we show that (1) the plan
 91 compliance problem is PSPACE-complete if no branching actions are allowed and (2) is
 92 EXPTIME-complete if branching actions are allowed. We also investigate the complexity
 93 of the reachability problem and the timed system compliance problems.
- 94 3. We present a formal semantics of our model based on logic, namely, on linear logic with
 95 definitions [36, 5, 4]. In particular, we provide an encoding and prove that there is a
 96 one-to-one correspondence between the plans and the (cut-free) focused proofs [2] of its

4 A Rewriting Framework and Logic for Activities Subject to Regulations

97 encoding.

98 Regarding contribution **1** described above, even in the case of balanced actions, we have to
99 deal with the problem that a plan can generate timestamps T of unbounded numeric values. In
100 particular, the state space is internally infinite since an *arbitrary* number of time advances can
101 occur (as illustrated at the beginning of Section 5). In our previous work [20] we were able to
102 solve a similar unboundedness problem caused by the presence of freshly created objects that
103 are called *nonces* in protocol security literature. However, the solution proposed in [20] is not
104 applicable to the problem of unboundedness of time. As a result, in this paper we have made
105 special precautions in our choice of a novel equivalence relation among states based on the
106 time differences of the timestamps of facts. This allows us to cover all plans of unbounded
107 length caused by uncontrolled time advances, with providing our upper bounds for the timed
108 collaborative systems (Theorem 5.6). We also show that our new technique introduced in this
109 paper can be combined with the technique introduced in [20] to solve the unboundedness for
110 both time and nonces in timed systems. In our experiments, we used this novel equivalence
111 relation among states.

112 The paper is organized as follows. Section 2 introduces the formal model for timed
113 collaborative systems called Timed Local State Transition Systems (*TLSTS*) as well as the
114 plan compliance problem described above. (In [32], *TLSTS*s were only mentioned, but not
115 formally introduced.) In Section 4 we give a formal semantics of our model based on focused
116 proofs of linear logic with definitions. Section 5 introduces an equivalence relation between
117 states of the system that allows us to handle the unboundedness of time with a finite space.
118 The machinery introduced in this section is used in Section 6 to demonstrate the decidability
119 of the plan compliance problem. Section 6 contains the complexity results mentioned above.
120 Section 7 we show that relaxing any of the main conditions on rules described above leads
121 to the undecidability of the reachability problem and thus the undecidability of the other
122 compliance problem described above. Finally in Sections 9 and 10 we discuss related and
123 future work.

124 This paper extends the conference paper [23] by providing in Section 4 a linear logic
125 semantics to our model based on multiset rewriting and time constraints and also by providing
126 full details for our complexity results for the plan compliance problem. In addition we
127 investigate the reachability problem and the system compliance problem, which were not dealt
128 in our conference paper [23]. We also show in Section 7 that the restrictions we impose on the
129 form of actions and time constraints in our systems are necessary for obtaining the decidability
130 of these problems. Relaxing any of those restrictions leads to undecidability. These results are
131 also novel with respect to our previous work [23].

132 2 Basic Definitions

133 At the lowest level, we have a first-order alphabet Σ that consists of a set of predicate symbols
134 P_1, P_2, \dots , function symbols f_1, f_2, \dots , constant symbols c_1, c_2, \dots , and variable symbols
135 x_1, x_2, \dots all with specific sorts (or types). The multi-sorted terms over the alphabet are
136 expressions formed by applying functions to arguments of the correct sort. Since terms may
137 contain variables, all variables must have associated sorts. A fact is an atomic predicate over
138 multi-sorted terms.

139 In order to accommodate the dimension of time in our model, we associate to each fact a
140 timestamp. *Timestamped facts* are of the form $P(t_1, \dots, t_n)@t$, where the number t is the
141 timestamp of the fact $P(t_1, \dots, t_n)$. Among the set of predicates, we distinguish the zero

142 arity predicate *Time*, which intuitively denotes the current global time of the system. For
 143 instance, the fact $Time@2$ denotes that the global time is 2. Here, we assume that timestamps
 144 are natural numbers. The intuitive meaning of a timestamp may depend on the system one is
 145 modeling. For instance, in our clinical investigations example, the timestamp associated to a
 146 fact could denote the time when a problem with a subject has been detected.

147 The *size of a fact*, P , denoted by $|P|$, is the total number of symbols it contains. We count
 148 one for each constant, variable, predicate, and function symbol, *e.g.*, $|P(f(x))| = 3$, and
 149 $|P(x, c, x)| = 4$. For our complexity results, we assume an upper bound on the size of facts, as
 150 in [15, 25, 20]. This means that for all facts, $P(t_1, \dots, t_n)@t$, the arity of predicate symbols,
 151 n , and the depth of terms, t_1, \dots, t_n , are bounded. However, we make no assumptions on the
 152 depth of timestamps, t , that is, the size of timestamps may be unbounded.

153 A *state*, or *configuration* of the system is a finite multiset, $Q_1@t_1, \dots, Q_n@t_n$, of *grounded*
 154 timestamped facts, *i.e.*, timestamped facts not containing variables. Configurations are assumed
 155 to contain exactly one occurrence of the predicate *Time*. We use W, X to denote the multiset
 156 resulting from the multiset union of W and X . For instance, the configuration

$$\{Time@5, Blood(id_1, scheduled)@7, Dose(id_1)@4, Status(id_1, normal)@5\}$$

157 denotes that that current time is 5, that the blood test for subject identified by id_1 should be
 158 taken on time 7, that the same subject took a dose of the drug at time 4, and his status is
 159 *normal*, *i.e.*, no problem has been detected.

160 For simplicity we often omit the word “timestamped” and just use the wording fact.

161 Following [25], we assume that the global configuration is partitioned into different local
 162 configurations each of which is accessible only to one agent. There is also a public configura-
 163 tion, which is accessible to all agents. As argued in [25], this separation allows one to specify
 164 systems for which it is important to know which facts are owned and can be manipulated
 165 by an agent of the system. Formally, this separation of the global configuration is done by
 166 partitioning the set of predicate symbols in the alphabet and it will be usually clear from
 167 the context. The time predicate *Time* is assumed to be public. For instance, in the above
 168 configuration all facts, except *Time*, belong to the health institution monitoring the subject
 169 id_1 .

170 **Time constraints** The time requirements of a system are specified by using *time con-*
 171 *straints*. Time constraints are arithmetic comparisons involving exactly two timestamps:

$$T_1 = T_2 \pm d, T_1 > T_2 \pm d, \text{ or } T_1 \geq T_2 \pm d, \quad (2.1)$$

172 where d is a natural number and T_1 and T_2 are time variables, which may be instantiated by
 173 the timestamps of any fact including the global time.

174 A concrete motivation for time constraints to be relative is that, as in physics, the rules of a
 175 collaboration are also not affected by time shifts. If we shift the timestamps of all facts by
 176 the same value, the same rules and conditions valid with respect to the original state are also
 177 valid with respect to the resulting state. If time constraints were not relative, however, then
 178 one would not be able to establish this important invariant. Indeed, as we show in Section 7,
 179 the reachability problem is undecidable for systems with non-relative time constraints.

180 2.1 Branching Actions and Plans

181 **Branching Actions** Actions work as multiset rewrite rules. As in [25, 20] we assume that
 182 each agent has a finite set of actions. However, we extend actions in two different ways:
 183 First, we add *guards to actions*; and second we allow actions to have a finite number of
 184 *non-deterministic effects*.

185 In their most general form, actions have the following form:

$$W \mid \mathcal{T} \longrightarrow_A [\exists \vec{x}_1. W_1] \oplus \cdots \oplus [\exists \vec{x}_n. W_n] \quad (2.2)$$

186 The subscript A is the name of the agent that owns this action. W is the pre-condition of
 187 this rule, while W_1, \dots, W_n are its post-conditions. All facts in W, W_1, \dots, W_n are public
 188 and/or belong to the agent A . \mathcal{T} is the guard of the action consisting of finitely many *time*
 189 *constraints* of the form shown in Equation 2.1. The existentially quantified variables specify
 190 the creation of fresh values, also known as nonces in protocol security literature.¹ Finally,
 191 if $n > 1$, then we classify the action as *branching*, otherwise, when $n = 1$, we classify the
 192 action as *non-branching*.

193 We say that a rule r of the form shown in Equation 2.2 creates a fact $F@T$, if $F@T$ does
 194 not appear in its pre-condition W , but appears in at least one of its post-conditions W_i .

195 With the exception of Section 7, we only consider in this paper systems with actions of the
 196 form shown in Equation 2.2 that may be of the following two types:

197 (*Time Tick Action*) The first one is the following action belonging to the special agent *clock*:

$$Time@T \mid \{\} \rightarrow_{clock} Time@(T + 1). \quad (2.3)$$

198 The above action does not have any constraints, which is specified by the empty set $\{\}$. It is
 199 the only action of the agent *clock* and is the *only action* that can change the global time.

200 (*Atomic Actions*) The second type of actions are those belonging to the remaining agents.
 201 We impose the following two conditions on actions depicted in Equation 2.2. Firstly, the
 202 global time $Time@T$ appears in the pre-condition, W , and in each of the post-conditions
 203 W_1, \dots, W_n exactly once. Secondly, if $Time@T$ is in the pre-condition W , then all facts
 204 created by the rule are of the form $P@(T + d)$, where d is a natural number, possibly zero.
 205 That is, all the facts created by this action have timestamps greater or equal to the global time.
 206 Notice that in this type of action the timestamp of *Time* does not change, that is, these actions
 207 are *instantaneous*. Also notice that, for example, the following action is not allowed

$$Time@T, R@T_1, P@T_2 \mid T_1 < T \longrightarrow_A Time@T, R@T_1, S@T_1.$$

208 This is because the timestamp of the created fact $S@T_1$ is not of the form $(T + d)$. That is,
 209 actions cannot create facts with arbitrary timestamps, instead they are only allowed to create
 210 facts whose timestamps are in the present or in the future, that is equal to or greater than the
 211 current time.

212 As we discuss in Sections 5 and 7, the two conditions on the actions belonging the agents
 213 different from the *clock* agent, discussed above, play an important role for the decidability of
 214 the system.

¹Fresh values are also often used in administrative processes, such as when a transaction number is issued. In particular, the transaction number has to be fresh. For a more detailed account for fresh values in administrative processes, see [19].




FIG. 1: A branching plan obtained by applying an action α of the form shown in Equation 2.2. Here σ is a ground substitution for α 's pre-condition W , while $W'_1\sigma, \dots, W'_n\sigma$ are ground instantiations of α 's post-conditions.

215 **Branching Plans** A *branching plan*, or simply *plan* is a tree whose nodes are configura-
 216 tions and whose edges are labeled with a pair consisting of an action and a number, $\langle \alpha, i \rangle$. As
 217 depicted in Figure 1, a plan is constructed by applying an action to one of its leaves. Formally,
 218 consider a branching action α of the form shown in Equation 2.2, that is, with pre-condition
 219 W and post-condition $W_1 \oplus \dots \oplus W_n$. We enumerate the post-conditions as W_1, \dots, W_n .
 220 When such an action is applied to a leaf of a plan labeled with W_I , the corresponding branch
 221 of the plan is extended by adding n leaves. The configuration labeling the i^{th} leaf is obtained
 222 by replacing α 's pre-condition, $W\sigma$, instantiated by a ground substitution σ in W_I by the
 223 corresponding post-condition of α , $W_i\sigma$, instantiated by the same substitution σ . The edge
 224 connecting W_I with i^{th} new leaf is labeled with $\langle \alpha, i \rangle$. In the process fresh values are created,
 225 replacing the existentially quantified variables, \vec{x}_i .

226 For example, let $\{Time@6, P(t_1)@1, Q(t_2)@4\}$ be a configuration appearing in a leaf of a
 227 plan \mathcal{P} . Then the following branching action is applicable:

$$Time@T, Q(Y)@T_1 \mid \{T > T_1 + 1\} \longrightarrow_A [\exists x. Time@T, R(Y, x)@T] \oplus [Time@T, S(Y)@T]$$

228 and it extends the plan \mathcal{P} creating the following two leaves $\{Time@6, P(t_1)@1, R(t_2, z)@6\}$
 229 and $\{Time@6, P(t_1)@1, S(t_2)@6\}$, where z is a fresh value.

230 If only non-branching actions are used, the plan has a single branch, *i.e.* the plan is simply a
 231 sequence of actions.

DEFINITION 2.1

232 A *timed local state transition system (TLSTS)* \mathcal{T} is a tuple $\langle \Sigma, I, R_{\mathcal{T}} \rangle$, where Σ is the alphabet
 233 of the language, I is a set of agents, such that $clock \in I$, and $R_{\mathcal{T}}$ is a finite set of actions
 234 owned by the agents in I of the two forms described above.

235 **Balanced Actions** We classify an action as *balanced* if its post-conditions, W_i , and the
 236 pre-condition, W , have the same number of facts. In Equation 2.2, this means that the number
 237 of facts in W and W_i are the same for all $1 \leq i \leq n$. We classify a *TLSTS* as *balanced* if all
 238 its actions are balanced.

239 For any plan P obtained from a balanced system, one can easily prove that *all configurations*
 240 *in P have the same number of facts*, namely the number of facts in P 's initial configuration.
 241 Intuitively, this means the number of facts that can be owned by an agent in the system is
 242 bounded by the number of facts in the initial configuration. In the remainder of this paper, we
 243 use the letter m to denote this number. Moreover, since we assume facts to have a bounded
 244 size, denoted using the letter k , the use of balanced actions imposes roughly a bound on the
 245 storage capacity of the agents in the system. In particular, any configuration in a plan obtained

8 A Rewriting Framework and Logic for Activities Subject to Regulations

246 from a balanced system, may have at most mk symbols. For more about balanced systems,
247 we point the reader to [25, 19].

248 As we further discuss in Section 6, the assumption that all actions in the systems are balanced
249 is crucial for showing that the reachability problem is in PSPACE. In fact, it was shown in
250 previous work [21] that this problem is undecidable if we allow actions to be un-balanced.

251 2.2 Planning Problems

252 In a collaboration, agents interact in order to achieve some common goal. However, since they
253 do not trust each other completely, they also want to avoid some critical situations. Often these
254 goals and critical situations mention time explicitly. For instance, in the clinical investigations
255 example discussed in the Introduction, the participants want to collect conclusive data without
256 violating regulations. Moreover, the sponsor should send a safety report to the FDA whenever
257 a serious and unexpected problem is detected within 15 days. Otherwise, the sponsor can be
258 severely penalized.

259 In order to formalize such aspects of a collaboration, we extend the notion of *initial, goal and*
260 *critical configurations* proposed in [25] by attaching a set of time constraints to configurations.
261 In particular, timed initial, goal and critical configurations have the following form:

$$\{Q_1@T_1, Q_2@T_2, \dots, Q_n@T_n\} \mid \mathcal{Y}$$

262 where \mathcal{Y} is a finite set of time constraints as shown in Eq. 2.1 such that its variables are in
263 T_1, T_2, \dots, T_n .

264 For instance, in the clinical investigations example, a possible goal configuration is the one
265 representing a situation when the data of a subject is collected in specified intervals for some
266 number of times. The following goal configuration specifies that the goal is to collect the data
267 of a subject 25 times in intervals of 28 days, but with a tolerance of 5 days:

$$\{Time@T, Data(Id, 1)@T_1, \dots, Data(Id, 25)@T_{25}\}$$

268 with the time constraints $T_i + 23 \leq T_{i+1} \leq T_i + 33$ and that $T > T_i$, for $1 \leq i \leq 25$.
269 Formally, any instantiation of the variables T_1, \dots, T_{25} that satisfies the set of constraints
270 above is considered a goal configuration.

271 Similarly, a configuration is critical for the participants of a clinical investigation when a
272 problem is detected at time T_1 , but no written report is sent to the FDA on time, *i.e.*, within 15
273 days after the problem is detected:

$$\{Detect(Id)@T_1, Report(Id)@T_2\} \mid \{T_2 > T_1 + 15\}.$$

274 Adding time constraints to configurations is not a restriction of the model. Quite the contrary,
275 time constraints provide a general mechanism to specify in a succinct fashion the set of goal
276 and critical configurations expressing time requirements.

277 For simplicity, we often omit the word "timed" in initial/goal/critical configurations regard-
278 less of time constraints being attached or not.

279 As in [19], we assume that the goal and critical configurations are closed with respect to
280 fresh values. That is, if a configuration C containing some nonces is a goal (respectively,
281 critical) state, then $C\sigma$ is also a goal (respectively, critical), where σ is a renaming of nonce
282 names. This assumption is sensible, as when defining critical and goal configurations, one
283 cannot specify the nonce names in advance, since these are freshly generated during the

284 execution of the process being modeled. The particular nonce name should not matter for
 285 classifying a configuration as critical or a goal configuration.

286 **Planning Problems** In [25, 21] three compliance problems were introduced in the setting
 287 without explicit time or branching (actions with non-deterministic effects). We now restate
 288 two if these problems in our setting with explicit time and branching.²

289 Given an initial configuration W_I and a finite set of goal and critical configurations, we call
 290 a *branching plan* \mathcal{P} *compliant* if it does not contain any critical configurations and moreover
 291 if all branches of \mathcal{P} lead from the initial configuration W_I to a goal configuration.

- 292 • (Timed plan compliance problem) Given a timed local state transition system \mathcal{T} , an initial
 293 configuration W consisting of timestamped facts and a finite, possibly empty, set of time
 294 constraints, a timed goal configuration Z , and a finite set of timed critical configurations,
 295 is there a compliant plan which leads from W to Z ?
- 296 • (Timed system compliance problem) Given a timed local state transition system \mathcal{T} , an
 297 initial configuration W consisting of timestamped facts and a finite, possibly empty,
 298 set of time constraints, a timed goal configuration Z , and a finite set of timed critical
 299 configurations, is no timed critical configuration reachable from W , and does there exist a
 300 plan leading from W to Z ?

301 In [21], the plan compliance problem without explicit time was called weak plan compliance.

302 Although the above problems are stated as decision problems, we prove more than just
 303 existence of a plan. Ideally, we are also able to generate a plan when there is a solution.
 304 Unfortunately, the number of actions in the plan may be very large, potentially increasing
 305 the complexity of the plan generation. For this reason we follow [25] and use the notion
 306 of “scheduling” a plan. However, since we are dealing with branching plans, whereas [25]
 307 considered non-branching plans, we need to agree how the nodes of a branching tree are
 308 enumerated. Therefore, we assume fixed a *tree traversal procedure*. It can be any traversal
 309 procedure, for instance, depth-first traversal procedures (pre, in-order, or post-order) or a
 310 breadth-first traversal procedure. Assuming such a tree traversal procedure, a scheduling
 311 algorithm takes an input i representing the node in the agreed traversal and outputs the i^{th}
 312 action of the plan, which extends this node.

DEFINITION 2.2

313 Assume pre-defined any tree traversal procedure. An algorithm is said to *schedule* a plan if it
 314 (1) finds a plan if one exists, and (2) on input i , if the plan contains at least i nodes, then it
 315 outputs the i^{th} action of the plan, otherwise it outputs *no*.

316 3 Implementing a TLSTS in Maude

317 The general-purpose computational tool Maude [9] provides all the machinery necessary to
 318 implement TLSTS specifications directly. As Maude is based on rewriting, the Maude code
 319 looks similar to the specification itself. We now illustrate this by using examples of how the
 320 encoding works.

²The third compliance problem, introduced as the *plan compliance problem* in [21], was called *semi-critical plan compliance problem* in [19] where it was observed that, for systems without explicit time, this problem is reducible to an instance of the plan compliance problem with a larger set of critical configurations. This set includes the set of *semi-critical configurations* from which it is possible to reach a critical state of a particular agent without the participation of this agent. The same reduction can be obtained for TLSTSes.

10 A Rewriting Framework and Logic for Activities Subject to Regulations

321 **Configurations** We start by specifying the signature of a *TLSTS*, *i.e.*, the set of constants
322 and predicate symbols. For instance, the code below specifies that the zero arity fact `time`
323 is of sort (or type) `Fact` and that `blood` is a binary fact whose argument is of sort `Id` and
324 `Result`.

```
325   op time : -> Fact .   op blood : Id Result -> Fact .
```

326 Other predicates of the sort `Fact` can be specified in a similar fashion.

327 We specify the operator `@` which attaches a natural number to facts as follows. It is used to
328 specify timestamped facts which are of sort `TFact`.

```
329   op @_ : Fact Nat -> TFact .
```

330 To encode configurations, we first specify that the sort of timestamped facts is a subset of
331 the sort configuration, denoted by the symbol `<`, that the empty set is a configuration, specified
332 by the operator `none`, and that the juxtaposition of two configurations is also a configuration.

```
   subsort TFact < Conf .
```

```
333   op none : -> Conf .
```

```
   op -- : Conf Conf -> Conf [assoc comm id:none] .
```

334 The last statement also specifies that configurations are multisets by attaching the keywords
335 `assoc` and `comm`, which specify that the operator constructing configurations is both asso-
336 ciative and commutative. Hence, when Maude checks whether an action (specified below) is
337 applicable, Maude will consider all possible permutations of elements until it finds a match
338 which satisfies the action's pre-condition as well as its guard. Finally, the keyword `id:none`
339 specifies that the constructor `none`, specifying the empty set, is the identity of an operator. It
340 is used to identify configurations, for example, the configurations below are identified

```
   none (time@2) none (blood(id1,positive)@3) and  
   (time@2) (blood(id1,positive)@3)).
```

341 **Timed Critical and Timed Goal Configurations** Timed critical and timed goal confi-
342 gurations are specified by *equational theories*. For instance, the following equational theory in
343 Maude specifies the critical configuration when the FDA is not notified within 7 days after
344 a serious and unexpected problem is detected. Here `Num` is the fresh value, *e.g.*, a number,
345 uniquely identifying a serious and unexpected event with subject identified by `Id`.

```
346   ceq critical((C:Conf) (time@T) (detected(Id,Num)@T1)  
   (fda(Id,no,Num)@T2)) = true if T > T1 + 7
```

347 Maude automatically replaces `critical(C)` with the boolean `true` if the configuration
348 `C` satisfies the condition specified by the equation above. Timed goal configurations are also
349 specified as equational theories in a similar way, only that we use the predicate `goal`, instead
350 of `critical` to specify goal configurations.

351 **Branching Actions and Searching for Compliant Plans** Whereas critical and goal
352 configurations are specified by using equational theories, actions are specified as rewrite rules
353 in Maude. To accommodate branching actions, we use three new operators `noPlan`, denoting
354 when a branching plan has no leaves, brackets used to mark a leaf of a plan, and `+` used to
355 construct the list of leaves of a branching plan. The leaves of a branching plan are of the sort
356 `Plan`.

```
   op noPlan : -> Plan .
```

```
357   op {-} : Conf -> Plan .
```

```
   op _+ : Plan Plan -> Plan [assoc id:noPlan] .
```

358 The operator $+$ is also used to specify the different outcomes of an action. For instance, the
359 following conditional rule specifies that there are two possible outcomes when a blood test
360 scheduled at time T_1 is carried out, namely, the blood test is positive or negative.
361 Moreover, the boolean conditions specifies that the test can only be carried out at the same day
362 when it was scheduled and if none of its outcomes is a critical configuration.

```
362 crl[blood]: { (C:Conf) (time@T) (blood(Id, scheduled)@T1) } =>
              { (C:Conf) (time@T) (blood(Id, positive)@T) } +
              { (C:Conf) (time@T) (blood(Id, negative)@T) }
363 if T1 = T ^
   not (critical((C:Conf) (time@T) (blood(Id, positive)@T))) ^
   not (critical((C:Conf) (time@T) (blood(Id, negative)@T)))
```

364 Formally, when this rule is applied then two different leaves are created, one for each
365 possible result. The remaining facts appearing in the configuration C are left untouched.

366 Notice that the definition of the $_{+}$ operator does not specify it to be commutative. However,
367 regarding the compliance problem that we are interested in (described in Section 2), changing
368 the order of the branches of a plan preserves its compliance as the resulting plan does not
369 reach any critical configuration and each of its leaves are goal configurations. Thus, we can
370 safely change the definition of $_{+}$ to also be commutative. As we demonstrate in Section 8,
371 this change reduces the number of possible states in average by a factor of 8.

372 As in the rule above, we allow a rule to be applied only if *all* its outcomes are *not* critical
373 configurations. For instance, the action that advances time (Eq. 2.3) is specified in Maude with
374 an extra condition allowing the time to be incremented only if the resulting configuration is
375 not critical:

```
376 crl[time]: { (C:Conf) (time@T) } => { (C:Conf) (time@(T+1)) }
   if not (critical((C:Conf) (time@(T+1))))
```

377 This means that it is not possible to reach a critical configuration when using the rules as
378 encoded above. Therefore, in order to search for a compliant plan, one does not need to care
379 whether a critical configuration is reached, as this is not possible, but only check whether there
380 is a plan from an initial configuration to a goal configuration obtained by using the actions as
381 mentioned above. Maude can automatically perform this search by using a command of the
382 following form:

```
383 search in MODULE_NAME : I =>+ P:Plan
   such that goals(P:Plan) = true .
```

384 where I is the initial configuration, $MODULE_NAME$ is the name of the Maude module contain-
385 ing all the rules of the *TLSTS*, and finally $goals$ is a boolean function (predicate) specified
386 by an equational theory that returns true when given $\{C_1\} + \dots + \{C_n\}$ of type $Plan$
387 only if $goal(C_i)$ evaluates to true for all $1 \leq i \leq n$.

388 It is often possible to demonstrate the non-interference of two actions, α and β , syntactically.
389 For instance, if there is no intersection between the facts modified by α and β , these actions
390 do not interfere between each other as they mention different parts of a configuration. The
391 following action specifying a vital sign test does not interfere with the action above specifying
392 a blood test:

```
393 crl[vital]: { (C:Conf) (time @ T) (vital(I, ID, false)@T1) } =>
              { (C:Conf) (time @ T) (vital(I, ID, true)@T) }
   if T1 = T ^
   not (critical((C:Conf) (time@T) (vital(I, ID, true)@T)))
```

394 This means that a compliant plan containing a sequence of actions $\alpha; \beta$ can be replaced
395 with another compliant plan where the order is inverted $\beta; \alpha$. In our example scenario, such

interleavings increase the number of states Maude must explore by a factor of 23. A better approach would be to merge these actions into a (big-step) action. For example, the big-step action obtained from the two actions above would specify the actions of performing the vital signs and blood test at the same time. For instance, one of its post-conditions specifies when the blood test is positive:

$$\{ (C:\text{Conf}) (\text{time}@\text{T}) (\text{vital}(\text{I}, \text{ID}, \text{true})@\text{T}) (\text{blood}(\text{Id}, \text{positive})@\text{T}) \}.$$

Finally, besides searching for plans, the same theory can also be used for monitoring CI executions. For instance, by using the equational theory specifying critical configurations, one can detect when a deviation has occurred and send alarms to the responsible agents. After a CI has been carried out, one could also use the actual plan carried out to study how CIs have been executed.

4 Formal Semantics using Linear Logic with Definitions

This Section provides a formal semantics for *TLSTSes* based on linear logic with definitions [36, 4]. In particular, we provide an encoding for *TLSTSes* such that given an initial configuration W and a *TLSTS*, then there is a one-to-one correspondence between the set of plans from W to a goal state Z and the set of (cut-free) focused proofs [2] of its encoding.

4.1 Focused Proof System for Linear Logic with Definitions

The focused proof system, LLF, for linear logic is depicted in Figure 2 and was introduced by Andreoli [2]. Focused proofs can be regarded as the normal form proofs for proof search. In order to formally introduce LLF, we first classify the connectives 1 , \otimes , \oplus , and \exists as positive and the remaining as negative. This distinction is natural as the introduction rules for the positive connectives are not-necessarily invertible, while the rules for the negative connectives are invertible. The same distinction, however, does not apply so naturally to literals and hence these are *arbitrarily* classified as positive or negative. Positive polarity literals and formulas whose main connective is positive are classified as positive formulas and the remaining as negative formulas.

As one can see from an inspection of LLF in Figure 2, there are two different sequents in LLF: those containing \uparrow which belong to the negative phase where only negative formulas are introduced, and those containing \downarrow which belong to the positive phase and only positive formulas are introduced. The decide rules D_1, D_2 , reaction rules $R \uparrow, R \downarrow$ and the bang introduction rule $!$ mark transition between positive and negative phases.

A key property of LLF is that it allows one to construct macro-rules that introduce synthetic connectives. For example, assume that the N_1, N_2, N_3 are all negative formulas. Then from the focusing discipline, there are only two possible ways to introduce the sequent $\vdash \Theta : \Gamma_1, \Gamma_2 \downarrow (N_1 \oplus N_2) \otimes N_3$:

$$\frac{\frac{\vdash \Theta : \Gamma_1 \uparrow N_1 \quad \vdash \Theta : \Gamma_2 \uparrow N_3}{\vdash \Theta : \Gamma_1, \Gamma_2 \downarrow (N_1 \oplus N_2) \otimes N_3}}{\vdash \Theta : \Gamma_1, \Gamma_2 \downarrow (N_1 \oplus N_2) \otimes N_3}} \quad \text{and} \quad \frac{\frac{\vdash \Theta : \Gamma_1 \uparrow N_2 \quad \vdash \Theta : \Gamma_2 \uparrow N_3}{\vdash \Theta : \Gamma_1, \Gamma_2 \downarrow (N_1 \oplus N_2) \otimes N_3}}{\vdash \Theta : \Gamma_1, \Gamma_2 \downarrow (N_1 \oplus N_2) \otimes N_3}}$$

Hence, the formula $(N_1 \oplus N_2) \otimes N_3$, under the focusing discipline, specifies such macro-rules which are obtained by applying the corresponding positive and a negative phase rules.

One can specify in a similar way a multiset rewrite rule r as a linear logical formula $F(r)$ in such a way that the macro-rule obtained by focusing on $F(r)$ corresponds *exactly* to the

Introduction Rules

$$\begin{array}{c}
 \frac{\vdash \Theta : \Gamma \uparrow L}{\vdash \Theta : \Gamma \uparrow L, \perp} [\perp] \quad \frac{\vdash \Theta : \Gamma \uparrow L, F, G}{\vdash \Theta : \Gamma \uparrow L, F \wp G} [\wp] \quad \frac{\vdash \Theta, F : \Gamma \uparrow L}{\vdash \Theta : \Gamma \uparrow L, ?F} [?] \\
 \\
 \frac{}{\vdash \Theta : \Gamma \uparrow L, \top} [\top] \quad \frac{\vdash \Theta : \Gamma \uparrow L, F \quad \vdash \Theta : \Gamma \uparrow L, G}{\vdash \Theta : \Gamma \uparrow L, F \& G} [\&] \quad \frac{\vdash \Theta : \Gamma \uparrow L, F[c/x]}{\vdash \Theta : \Gamma \uparrow L, \forall x F} [\forall] \\
 \\
 \frac{}{\vdash \Theta : \downarrow 1} [1] \quad \frac{\vdash \Theta : \Gamma \downarrow F \quad \vdash \Theta : \Gamma' \downarrow G}{\vdash \Theta : \Gamma, \Gamma' \downarrow F \otimes G} [\otimes] \quad \frac{\vdash \Theta : \uparrow F}{\vdash \Theta : \downarrow !F} [!] \\
 \\
 \frac{\vdash \Theta : \Gamma \downarrow F}{\vdash \Theta : \Gamma \downarrow F \oplus G} [\oplus_l] \quad \frac{\vdash \Theta : \Gamma \downarrow G}{\vdash \Theta : \Gamma \downarrow F \oplus G} [\oplus_r] \quad \frac{\vdash \Theta : \Gamma \downarrow F[t/x]}{\vdash \Theta : \Gamma \downarrow \exists x F} [\exists]
 \end{array}$$

Identity, Reaction, and Decide rules

$$\begin{array}{c}
 \frac{}{\vdash \Theta : A_p^\perp \downarrow A_p} [I_1] \quad \frac{}{\vdash \Theta, A_p^\perp : \downarrow A_p} [I_2] \quad \frac{\vdash \Theta : \Gamma, S \uparrow L}{\vdash \Theta : \Gamma \uparrow L, S} [R \uparrow] \\
 \\
 \frac{\vdash \Theta : \Gamma \downarrow P}{\vdash \Theta : \Gamma, P \uparrow} [D_1] \quad \frac{\vdash \Theta, P : \Gamma \downarrow P}{\vdash \Theta, P : \Gamma \uparrow} [D_2] \quad \frac{\vdash \Theta : \Gamma \uparrow N}{\vdash \Theta : \Gamma \downarrow N} [R \downarrow]
 \end{array}$$

FIG. 2: The focused proof system, LLF, for linear logic [2]. Here, L is a list of formulas, Θ is a multiset of formulas, Γ is a multiset of literals and positive formulas, A_p is a positive literal, N is a negative formula, P is not a negative literal, and S is a positive formula or a negated atom.

434 operational semantics of the rewrite rule r . But in order to specify in the same way the
 435 semantics of time constraints of *TLSTSes*, we need more machinery, namely definitions [36,
 436 4]. A *definition* is a finite set of *clauses* which are written as $\forall \vec{x}[P(\vec{x}) \triangleq B]$: here P is a
 437 predicate and every free variable of B (the *body* of the clause) is contained in the list \vec{x} . The
 438 symbol \triangleq is not a logical connective but is used to indicate a definitional clause. We consider
 439 that every defined predicate occurs at the head of exactly one clause. Introduction rules for
 440 definitions are shown below, where a definition can be unfolded on both the positive phase
 441 and the negative phase:

$$\frac{\vdash \Theta : \Gamma \downarrow B\theta}{\vdash \Theta : \Gamma \downarrow P(\vec{c})} [def \downarrow] \quad \frac{\vdash \Theta : \Gamma \uparrow L, B\theta}{\vdash \Theta : \Gamma \uparrow L, P(\vec{c})} [def \uparrow]$$

442 The proviso for both of these rules is: $\forall \vec{x}[P(\vec{x}) \triangleq B]$ is a definition clause and θ is the
 443 substitution that maps the variables \vec{x} to the terms \vec{c} , respectively. Thus, in either phase of
 444 focusing, if a defined atom is encountered, it is simply replaced by its definition and the proof
 445 search phase does not change.

446 We also include the rules for equality shown below:

$$\frac{\{(\vdash \Theta : \Gamma \uparrow L)\theta \mid \theta \in CSU(r, s)\}}{\vdash \Theta : \Gamma \uparrow L, r \neq s} [\neq_r] \quad \frac{}{\vdash \Theta : \cdot \downarrow r = r} [=_r]$$

447 where $CSU(s, r)$ denotes the complete set of unifiers of two terms. Since we are dealing with
 448 first-order logic terms, this set either contains one unifier, the most general unifier, or it is

empty when the terms r and s are not unifiable. Notice that right equality introduction rule behaves exactly as the rule [1]. The proof theory of inference rules such as these is well studied (see, for example, [5, 28, 4]). Linear logic with definitions admits cut-elimination of LLF with definitions [28]³ and the focusing discipline used above was shown to be complete [4]. This paper will only need the $=_r$ rule.

4.2 Encoding TLSTSes in Linear Logic with Definitions

Encoding Arithmetic Conditions We show how to express the semantics of *TLSTSes* as search for cut-free focused linear logic proof with definitions. In particular, we use the following definitions to specify, for example, the arithmetic operations of \leq , $<$ and $+$ that appear in constraints:

$$\begin{aligned}
 x \leq y & \triangleq [x = zr] \oplus [\exists x'y'.(x = s(x')) \otimes (y = s(y')) \otimes (x' \leq y')]. \\
 x < y & \triangleq [\exists y'.(x = zr) \otimes (y = s(y'))] \oplus [\exists x'y'.(x = s(x')) \otimes (y = s(y')) \otimes (x' \leq y')]. \\
 Plus(x, y, z) & \triangleq [(x = zr \otimes y = z)] \oplus [\exists x'z'.((x = s(x')) \otimes (z = s(z'))) \otimes Plus(x', y, z')].
 \end{aligned}$$

where natural numbers are expressed by using the successor function s and the constant zr denoting the natural number zero. For instance, the definition for \leq contains two disjuncts: the left disjunct specifies the base case when the value of x is zero, and the right disjunct the inductive case, where both x and y are the successors of two numbers x' and y' such that $x' \leq y'$. The other arithmetic operations can be specified in a symmetric way.

As observed in [33], the definitions above can be used to compute an arithmetic operation in a single focused step. This is because the body of all the definitions above is positive. Therefore, once one focuses on one of the atoms defined above, one does not lose focus anymore and hence a proof consists necessarily of a single positive phase. For example, if we focus on the atom $s(zr) \leq s(s(zr))$ one obtains the following derivation:

$$\frac{\frac{\frac{\overline{\vdash \Theta : \cdot \Downarrow s(zr) = s(zr)} \quad [\text{=}r] \quad \overline{\vdash \Theta : \cdot \Downarrow s(s(zr)) = s(s(zr))} \quad [\text{=}r] \quad \vdash \Theta : \cdot \Downarrow zr \leq s(zr)}{\vdash \Theta : \cdot \Downarrow s(zr) = s(zr) \otimes s(s(zr)) = s(s(zr)) \otimes zr \leq s(zr)} \quad [2 \times \otimes]}{\vdash \Theta : \cdot \Downarrow \exists x'y' s(zr) = s(x') \otimes s(s(zr)) = s(y') \otimes x' \leq y'} \quad [2 \times \exists]}{\frac{\vdash \Theta : \cdot \Downarrow [s(zr) = zr] \oplus [\exists x'y' s(zr) = s(x') \otimes s(s(zr)) = s(y') \otimes x' \leq y']}{\vdash \Theta : \cdot \Downarrow s(zr) \leq s(s(zr))} \quad [\oplus_2]}{\vdash \Theta : \cdot \Downarrow s(zr) \leq s(s(zr))} \quad [\text{def}\Downarrow]$$

At the open branch, the definition for the atom $zr \leq s(zr)$ is necessarily unfolded and the left disjunct of its definition is used to finish the proof. Notice that under the focusing discipline there is no other way to introduce a sequent focused on the atom $s(zr) \leq s(s(zr))$. If, for example, one attempts to prove the sequent by choosing instead the left disjunct of its body definition, one would fail since it is not possible to introduce a sequent focused on the equality $s(zr) = zr$. For a similar reason, to obtain a proof, one has to instantiate the variables x'

³Technically it was shown that cut-elimination works when definitions satisfy certain conditions which are out of scope of this paper. The definitions that we need here fall under this fragment.

475 and y' with zr and $s(zr)$, respectively. Otherwise, it is not possible to introduce the resulting
 476 equalities.

477 **Encoding of Timestamped Facts and Constraints** Using above definitions, we encode
 478 a constraint of the form $T_1 \circ T_2 + d$ as a logical formula

$$[Plus(T_2, \ulcorner d \urcorner, T_2')] \otimes [T_1 \circ T_2'],$$

479 where $\circ \in \{>, \geq, =, \leq, <\}$ and $\ulcorner d \urcorner$ is the term corresponding to the natural number d . The
 480 encoding contains the constants s and zr . For instance, the natural number 2 is translated
 481 into the term $s(s(zr))$. Notice as well that this formula has only positive connectives and as
 482 illustrated above, once it is focused on, focusing is never lost. Therefore, we can use them
 483 to check in one positive phase whether a constraint is satisfied. If C is a constraint then we
 484 denote $\ulcorner C \urcorner$ as the logical formula obtained from C .

485 To encode a timestamped fact with predicate name P in linear logic, we use a new predicate
 486 name P' with arity increased by one. The encoding $\ulcorner P(\vec{c})@t \urcorner$ is the formula $P'(\vec{c}, t)$. We
 487 extend the definition of $\ulcorner \cdot \urcorner$ for constraints, natural numbers, and timestamped facts to multiset
 488 as usual.

489 The encoding of a configuration will be placed in the linear context of sequents, namely in
 490 the context Γ of the sequents $\vdash \Theta : \Gamma \uparrow$. As we show below, the encoding of rewrite rules will
 491 be placed in the classical context, that is, the context Θ . This is because rewrite rules can be
 492 used any number of times.

493 **Encoding of Actions** To encode an action of the form

$$W \mid \mathcal{Y} \rightarrow_A \exists \vec{t}_1. W_1 \oplus \cdots \oplus \vec{t}_n. W_n$$

494 in linear logic, we first need to specify the timestamps of the form $T + d_i$ appearing in the
 495 post-conditions W_j s. For this, we construct two sets W_j^f and W_j^c from W' : for each fact
 496 of the form $Q_i@T + d_i$ in W' we add $Q_i@T_i$ in W_j^f , where T_i is a new variable, and the
 497 formula $Plus(T, \ulcorner d_i \urcorner, T_i)$ to W_j^c ; and for each fact of the form $Q_i@(T)$ in W' we add the
 498 same fact to W_j^f and no formula in W_j^c . Intuitively, the set W_j^c specifies the values for the
 499 new time variables used in W_j^f to be the same as specified in the original timestamps in W .
 500 One could regard the set W_j^c as a set of constraints to the new time variables introduced. For
 501 instance, the post-condition of the following action,

$$Time@T, P(x)@T_1, Q(y)@T_2 \mid \{T_2 > T_1 + 1\} \rightarrow_A \exists u. Time@T, P(u)@(T+2), R(x)@T,$$

502 returns the sets $\{Time@T, P(u)@(T_2'), R(x)@T\}$ and $\{Plus(T, s(s(zr)), T_2')\}$.

503 Now, we are ready to encode actions in linear logic: an action of the form $W \mid \mathcal{Y} \rightarrow_A$
 504 $\exists \vec{t}_1. W_1 \oplus \cdots \oplus \vec{t}_n. W_n$ is encoded as the linear logic formula

$$F = \forall \vec{x} \left[\bigotimes \ulcorner W \urcorner \otimes q_A \otimes \bigotimes \ulcorner \mathcal{Y} \urcorner \otimes \bigotimes_{j=1}^{j=n} W_j^c \multimap \bigoplus_{j=1}^{j=n} \left[\exists \vec{t}. \bigotimes \ulcorner W_j^f \urcorner \otimes q_A \right] \right],$$

505 where \vec{x} are the free variables appearing in the rule together with all the new variables
 506 introduced by the translation $\ulcorner \cdot \urcorner$ and in the set W_j^c . Also, the atomic formula q_A is used only

16 *A Rewriting Framework and Logic for Activities Subject to Regulations*

507 to mark that this action belongs to agent A . Moreover, the encoding of a set of transition rules
 508 $\ulcorner R_{\mathcal{T}} \urcorner$ is the set with the encoding of all the transition rules in $R_{\mathcal{T}}$, and the set of propositions
 509 used to mark a rule to an agent is defined as $Q_I = \{q_A : A \in I\}$. Intuitively, the encodings of
 510 actions are placed in the unbounded context in the left-hand-side of a sequent. However, since
 511 we are using a one-sided proof system, we use its negation in the one-sided LLF system with
 512 definitions:

$$F^\perp \equiv \exists \vec{x} \left[\left(\bigotimes^{\ulcorner W^\neg \urcorner} \otimes q_A \otimes \bigotimes^{\ulcorner \mathcal{Y}^\neg \urcorner} \otimes \bigotimes_{j=1}^{j=n} W_j^c \right) \otimes \bigotimes_{j=1}^{j=n} \left[(\forall t^\neg \mathcal{Y}^{\ulcorner W_j^{f \neg \perp} \urcorner} \wp q_A^\perp) \right] \right].$$

513 Assume now that all atomic formulas have positive polarity, and consequently their negation
 514 negative polarity. The focused derivation introducing F^\perp necessarily has to be of the form
 515 below. Recall that the encodings of rewrite rules are in the classical context Θ , thus $F^\perp \in \Theta$.

$$\frac{\frac{\frac{\vdash \Theta : \Delta \Downarrow \bigotimes^{\ulcorner W^\neg \urcorner} \otimes \bigotimes^{\ulcorner \mathcal{Y}^\neg \urcorner} \otimes \bigotimes_{j=1}^{j=n} W_j^c \otimes q_A \quad \vdash \Theta : \Gamma \Downarrow \bigotimes_{j=1}^{j=n} \left[(\forall t^\neg \mathcal{Y}^{\ulcorner W_j^{f \neg \perp} \urcorner} \wp q_A^\perp) \right]}{\vdash \Theta : \Gamma, \Delta \Downarrow \left(\bigotimes^{\ulcorner W^\neg \urcorner} \otimes q_A \otimes \bigotimes^{\ulcorner \mathcal{Y}^\neg \urcorner} \otimes \bigotimes_{j=1}^{j=n} W_j^c \right) \otimes \bigotimes_{j=1}^{j=n} \left[(\forall t^\neg \mathcal{Y}^{\ulcorner W_j^{f \neg \perp} \urcorner} \wp q_A^\perp) \right]} [\otimes]}{\vdash \Theta : \Gamma, \Delta \Downarrow F^\perp} [n \times \exists]}{\vdash \Theta : \Gamma, \Delta \uparrow \cdot} [D_2]$$

516 Since \otimes is a positive connective, the left-premise is necessarily introduced by a completely
 517 positive phase introducing all tensors in $\bigotimes^{\ulcorner W^\neg \urcorner} \otimes q_A \otimes \bigotimes^{\ulcorner \mathcal{Y}^\neg \urcorner} \otimes \bigotimes_{j=1}^{j=n} W_j^c$ until one only
 518 focuses on atomic formulas. There are then two types of atomic formulas: the first type are
 519 atoms that have a definition, such as *Plus*, and those that do not have a definition, such as
 520 q_A . When one of the former is focused on, the focusing discipline forces its definition to
 521 be opened, thus computing the values of the timestamps of the facts in the post-conditions,
 522 specified in W_j^c . Moreover, as discussed above, these are proved without using any formulas
 523 from Δ . That is, the sequent below is proved in a single positive phase:

$$\vdash \Theta : \cdot \Downarrow \bigotimes_{j=1}^{j=n} W_j^c$$

524 The same happens when checking whether the rule guard is satisfied or not. In particular, the
 525 following sequent should be proved in a single positive phase

$$\vdash \Theta : \cdot \Downarrow \bigotimes^{\ulcorner \mathcal{Y}^\neg \urcorner}$$

526 and is provable if and only if the guard \mathcal{Y} is satisfied.

527 The second type of atoms are those that do not have definitions and appear in $\ulcorner W^\neg \urcorner$ and
 528 the fact q_A . Since these are assumed to have positive polarity, the only applicable rule when
 529 these are focused on is an initial rule. This forces Δ to be exactly the negation of the facts
 530 in $\ulcorner W^\neg \urcorner$ union the fact q_A^\perp . In contrast, for the right-premise of the derivation above, since
 531 \forall and \wp are negative connectives, the right-premise is necessarily introduced by a negative
 532 phase introducing these connectives. Hence, the macro-rule introducing an encoding of the
 533 transition rule is necessarily of the form, which corresponds to the one in Figure 1:

$$\frac{\frac{\vdash \Theta : \Gamma, q_A^\perp, \ulcorner W_1^{f \neg \perp} \urcorner \sigma \uparrow \cdot \quad \dots \quad \vdash \Theta : \Gamma, q_A^\perp, \ulcorner W_n^{f \neg \perp} \urcorner \sigma \uparrow \cdot}{\vdash \Theta : \Gamma, q_A^\perp, \ulcorner W^\neg \urcorner \sigma \uparrow \cdot}}$$

534 Notice that if the pre-condition or the constraints in \mathcal{Y} of the action are not satisfied, then
 535 there is no focused proof which focuses on the encoding of this transition. This handles
 536 the inductive case of the adequacy result of our encoding of *TLSTS* in Linear Logic with
 537 Definitions (Theorem 4.1).

538 **Encoding of Partial Goal** The base case of our adequacy result consists in checking if a
 539 partial goal is reached. This is specified in a similar way as before. Let the set of facts Z and
 540 the set of time constraints \mathcal{Y} constitute a goal configuration G . To check whether this goal
 541 configuration is reached, we encode G , written $\ulcorner G \urcorner$ as follows:

$$\exists \vec{x}. \bigotimes \ulcorner Z \urcorner \otimes \ulcorner \mathcal{Y} \urcorner \otimes \top,$$

542 where \vec{x} is the set of time variables appearing in the goal configuration. This formula is
 543 necessarily introduced by the following focused derivation:

$$\frac{\frac{\frac{\frac{\overline{\vdash \Theta : \Delta \Downarrow \bigotimes \ulcorner Z \urcorner}}{\vdash \Theta : \Gamma, \Delta \Downarrow \bigotimes \ulcorner Z \urcorner \otimes \ulcorner \mathcal{Y} \urcorner \otimes \top} [n \times \exists]}{\vdash \Theta : \Gamma, \Delta \Downarrow \exists \vec{x}. \bigotimes \ulcorner Z \urcorner \otimes \ulcorner \mathcal{Y} \urcorner \otimes \top} [D_2]}{\vdash \Theta : \Gamma, \Delta \Uparrow} [2 \times \otimes]}{\vdash \Theta : \Gamma \Downarrow \top} [R \Downarrow, \top]} [D_2]$$

544 As before, since all atoms are assigned with positive polarity, the focusing discipline forces
 545 that Δ contains exactly the negation of the facts appearing in $\ulcorner Z \urcorner$, that all constraints in
 546 \mathcal{Y} are satisfied, and that Γ contains the remaining facts in the sequent. That is, the current
 547 configuration is a goal configuration.

548 Given the discussion above, we prove the following connection between linear logic with
 549 definitions and reachability using *TLSTS* by induction on the height of derivation trees and
 550 on the height/length of branching plans.

THEOREM 4.1

Let $\mathcal{T} = \langle \Sigma, I, R_{\mathcal{T}} \rangle$ be a timed local transition system. Let W be an initial configuration and
 G be a goal configuration under the signature Σ . Then the sequent

$$\vdash \ulcorner R_{\mathcal{T}} \urcorner^{\perp} : Q_I^{\perp}, \ulcorner W \urcorner, \ulcorner G \urcorner \uparrow.$$

551 is provable in linear logic with definitions where $\ulcorner X \urcorner$ is the encoding as defined above iff
 552 there is a branching plan whose root is W and whose leaves contain G .

553 In fact, the adequacy we get is stronger than what is stated by the result above. The adequacy
 554 is on the level of derivations [34]. That is, proof search in the linear logic encoding corresponds
 555 exactly to search using the encoded *TLSTS*. However, we must also notice that our encoding
 556 only deals with reachability and not with the Planning Problem as we do not check whether a
 557 state is critical. But, one can check whether a critical state C is reachable from an initial state
 558 by specifying the reachability goal G to be the critical state C .

559 **Remark** The representation of *TLSTS* in Maude described in [32] serves as an exe-
 560 cutable specification and comes with tools for simulation, reachability analysis and model
 561 checking. The Linear Logic semantics of *TLSTS* provides a different and complimentary tool
 562 for reasoning both about inference strategies and about systems specified in the formalism.
 563 Interestingly there is a close correspondence between LL derivations and rewriting logic
 564 derivations for *TLSTS* specifications (and more generally for multiset rewrite systems) [18].

565 **5 Dealing with the Unboundedness of Time**

566 Comparing our *timed* collaborative models introduced here with the results on the *untimed*
 567 collaborative systems in our previous work [19], we meet with a number of the crucial
 568 difficulties. In the case of planning problems for the untimed systems with balanced actions,
 569 we are dealing with a *finite* (though huge) state space. Here the state space is *internally infinite*,
 570 since an arbitrary number of time advances is allowed in principle. For a straightforward
 571 example, consider a plan where time is eagerly advanced. That is, consider a plan with a single
 572 branch where time advances constantly:

$$Time@0, W \xrightarrow{clock} Time@1, W \xrightarrow{clock} Time@2, W \xrightarrow{clock} \dots$$

573 Since there are no bounds on the length nor depth of plans, the final value of the global time
 574 cannot be bounded in advance.

575 This section describes how to overcome the above problem by proposing an equivalence
 576 relation between configurations. The key idea is that since time constraints are relative, that
 577 is, they involve the difference of two timestamps, we do not need to keep track of the actual
 578 values of timestamps, in order to determine whether our time constraints are satisfied or not.

579 **Truncated time differences** In particular, we will store the time differences among
 580 the facts, but truncated by an upper bound. Formally, assume D_{max} be an upper bound on
 581 the numbers appearing explicitly in a given planning problem with the model \mathcal{T} - that is,
 582 the numbers in the actions and time constraints in \mathcal{T} , and in the initial, goal and critical
 583 configurations, for instance, the number d in Eq. 2.1. Then the *truncated time difference* of
 584 two timed facts $P@T_1$ and $Q@T_2$ with $T_1 \leq T_2$, denoted by $\delta_{P,Q}$, is defined as follows:

$$\delta_{P,Q} = \begin{cases} T_2 - T_1, & \text{provided } T_2 - T_1 \leq D_{max} \\ \infty, & \text{otherwise .} \end{cases}$$

585 Intuitively, we can truncate time differences without sacrificing soundness nor completeness
 586 because time constraints are relative as defined in Eq. 2.1. Hence, if the time difference of two
 587 facts is greater than the upper bound D_{max} , then it does not really matter how much greater it
 588 is, but just that it is greater. For instance, consider the time constraint $t_1 \geq t_2 + d$ involving
 589 the timestamps of the facts $P@t_1$ and $Q@t_2$. If $\delta_{Q,P} = \infty$, this time constraint is necessarily
 590 satisfied.

591 **Equivalence between configurations** We use the notion of truncated time differences
 592 introduced above to formalize the following equivalence relation among configurations.

DEFINITION 5.1

593 Given a planning problem with the *TLSTS* \mathcal{T} , let D_{max} be an upper bound on the numeric
 594 values appearing in \mathcal{T} and in the initial, goal and critical configurations. Let

$$\mathcal{S} = Q_1@T_1, Q_2@T_2, \dots, Q_m@T_m \quad \text{and} \quad \tilde{\mathcal{S}} = Q_1@\tilde{T}_1, Q_2@\tilde{T}_2, \dots, Q_m@\tilde{T}_m$$

595 be two configurations written in canonical way where the two sequences of timestamps
 596 T_1, \dots, T_m and $\tilde{T}_1, \dots, \tilde{T}_m$ are non-decreasing. (For the case of equal timestamps, we sort
 597 the facts in alphabetical order, if necessary.) Then \mathcal{S} and $\tilde{\mathcal{S}}$ are equivalent if for any $1 \leq i < m$
 598 either of the following holds:

$$T_{i+1} - T_i = \tilde{T}_{i+1} - \tilde{T}_i \leq D_{max} \quad \text{or both} \quad T_{i+1} - T_i > D_{max} \quad \text{and} \quad \tilde{T}_{i+1} - \tilde{T}_i > D_{max}.$$

599 In order to illustrate the above equivalence, assume that $D_{max} = 3$ and consider the
600 following two configurations:

$$\{R@3, P@4, Time@11, Q@12, S@14\} \text{ and } \{R@0, P@1, Time@6, Q@7, S@9\}.$$

601 According to the above definition, these configurations are equivalent since their truncated
602 time differences are the same. This can be observed by checking their *canonical* representation,
603 called δ -representation defined below.

DEFINITION 5.2

604 Let $\mathcal{S} = Q_1@T_1, Q_2@T_2, \dots, Q_m@T_m$ be a configuration written in canonical way where
605 the sequence of timestamps T_1, \dots, T_m is non-decreasing (for the case of equal timestamps,
606 we sort the facts in alphabetical order, if necessary) and let D_{max} be an upper bound in a
607 planning problem (as per Definition 5.1). The δ -representation of configuration of \mathcal{S} , denoted
608 by $\delta_{\mathcal{S}}$, is the tuple

$$\langle Q_1, \delta_{Q_1, Q_2}, Q_2, \delta_{Q_2, Q_3}, Q_3, \dots, Q_i, \delta_{Q_i, Q_{i+1}}, Q_{i+1}, \dots, Q_{m-1}, \delta_{Q_{m-1}, Q_m}, Q_m \rangle.$$

609 A δ -representation is constructed from a given configuration by sorting its facts according
610 to their timestamps and sorting facts in alphabetical order as tie-breaker. Then we compute
611 the time difference among two consequent facts, $\delta_{Q_i, Q_{i+1}}$. For instance, both configurations
612 given above have the following δ -representation:

$$\langle R, 1, P, \infty, Time, 1, Q, 2, S \rangle.$$

613 Here a value appearing between two facts, Q_i and Q_{i+1} , is the truncated time difference of
614 the corresponding facts, $\delta_{Q_i, Q_{i+1}}$, e.g., $\delta_{R, P} = 1$ and $\delta_{P, Time} = \infty$. It is also easy to see that
615 from the tuple above, one can compute the remaining truncated time differences. For instance,
616 $\delta_{Time, S} = 3$, since $1 + 2 = 3$, while $\delta_{R, Q} = \infty$, since $1 + \infty + 1 = \infty$.

617 We now formalize the intuition described above that using time differences that are truncated
618 by an upper bound instead of actual timestamps, we are able to determine whether a time
619 constraint is satisfied or not.

LEMMA 5.3

620 Let S and \tilde{S} be two equivalent configurations from Definition 5.1.

621 $\mathcal{S} = Q_1@T_1, Q_2@T_2, \dots, Q_n@T_n$ and $\tilde{\mathcal{S}} = Q_1@\tilde{T}_1, Q_2@\tilde{T}_2, \dots, Q_n@\tilde{T}_n$.
622 Then the following holds for all i and j such that $i > j$, and for all $a \leq D_{max}$:

$$\begin{aligned} T_i - T_j = a & \quad \text{if and only if} \quad \tilde{T}_i - \tilde{T}_j = a \\ T_i - T_j < a & \quad \text{if and only if} \quad \tilde{T}_i - \tilde{T}_j < a \\ T_i - T_j > a & \quad \text{if and only if} \quad \tilde{T}_i - \tilde{T}_j > a \end{aligned}$$

624 PROOF. The only interesting case is the last one, which can be proved by using the fact
625 that $a \leq D_{max}$ and that S and \tilde{S} are equivalent. Hence, $T_i - T_j > D_{max} > a$ is true if
626 and only if $\tilde{T}_i - \tilde{T}_j > D_{max} > a$ is true, and $D_{max} \geq T_i - T_j > a$ is true if and only if
627 $D_{max} \geq \tilde{T}_i - \tilde{T}_j > a$, since $T_i - T_j = \tilde{T}_i - \tilde{T}_j$. ■

628 Following Lemma 5.3, we say that a δ -representation Δ satisfies a constraint if a configura-
629 tion W , such that $\delta_W = \Delta$, satisfies that constraint.

630 **Handling time advances and action applications** Our next task is to show that our
631 equivalence relation using truncated time differences is well-defined with respect to actions.

632 That is, we show that actions preserve the equivalence among configurations. This will allow
633 us to represent plans using δ -representations only.

634 We extend action application to δ -representations. It follows from the Lemma 5.3 that the
635 same action is applicable in configurations with the same δ -representation. We, therefore,
636 say that an *action is applicable in a δ -representation Δ* if the same action is applicable in
637 a configuration W , such that $\delta_W = \Delta$. That is, any action a that is applicable in some
638 configuration S is applicable in its δ -representation δ_S , and the resulting δ -representation, δ'_S ,
639 is the δ -representation of S' , where $S \rightarrow_a S'$:

$$\begin{array}{ccc} \delta_S & \rightarrow_a & \delta_{S'} \\ \wr & & \wr \\ S & \rightarrow_a & S' \end{array} \quad (5.1)$$

640 This is well defined if it is independent of the choice of configurations. Recall that there are
641 two types of actions, namely time advances and instantaneous actions that belong to agents.
642 Time advances only change the timestamp denoting the global time while the rest of the
643 configuration remains unchanged. Therefore, when we advance time in a δ -representation,
644 the position of *Time* and the truncated time differences involving *Time* need to be updated.
645 Depending on concrete values of time differences, the fact *Time* may move to the right.

646 For example, for $D_{max} = 5$ and the configuration $\{R@0, P@1, Time@3, Q@5, S@7\}$
647 with the time advance action $Time@T \rightarrow_{clock} Time@(T + 1)$ we get

$$\{R@0, P@1, Time@3, Q@5, S@7\} \rightarrow_{clock} \{R@0, P@1, Time@4, Q@5, S@7\}$$

648

$$i.e. \quad \langle R, 1, P, 2, Time, 2, Q, 2, S \rangle \rightarrow_{clock} \langle R, 1, P, 3, Time, 1, Q, 2, S \rangle .$$

649 With another application of time tick action we then get:

$$\langle R, 1, P, 3, Time, 1, Q, 2, S \rangle \rightarrow_{clock} \langle R, 1, P, \infty, Q, 0, Time, 2, S \rangle .$$

Generally, the time advance action $Time@T \rightarrow_{clock} Time@(T + 1)$ applied to

$$\Delta = \langle Q_1, \delta_1, \dots, Q_{i-1}, \delta_{i-1}, Time, \delta_i, Q_{i+1}, \delta_{i+1}, \dots, \delta_{m-1}, Q_m \rangle$$

results in the following δ -representation Δ' , alphabetically sorted whenever truncated time differences are equal to 0:

$$\begin{aligned} & \langle Q_1, \delta_1, \dots, Q_{i-1}, [\delta_{i-1} + 1], Time, \delta_i - 1, Q_{i+1}, \delta_{i+1}, \dots, \delta_{m-1}, Q_m \rangle, \quad \text{if } \delta_i \geq 1 \\ & \langle Q_1, \delta_1, \dots, Q_{i-1}, \delta_{i-1}, Q_{i+1}, \dots, Q_{i+l}, [\delta_{i+l} + 1], Time, \delta_{i+l} - 1, \dots, \delta_{m-1}, Q_m \rangle, \\ & \quad \text{if } \delta_i = \delta_{i+l-1} = 0, \delta_{i+l} > 0 \end{aligned}$$

650 where $[d]$ denotes d , for $d < D_{max}$, and denotes ∞ otherwise.

651 In case $\Delta = \langle Q_1, \delta_1, \dots, \delta_{m-1}, Time \rangle$, then $\Delta = \langle Q_1, \delta_1, \dots, [\delta_{m-1} + 1], Time \rangle$.

For the application of instantaneous actions recall that the fact $Time@T$ remains unchanged, while some facts from the pre-condition of the action are replaced with other facts whose timestamps are of the form $T + d$. We modify the δ -representation in the following way. We first remove the facts that appear in the pre-condition of the action and not in its post-condition. Then we insert the new facts from the post-condition, positioning them on the basis of their time difference to the fact *Time*, and alphabetically if necessary. Finally, we fill

in the new time differences. This is best explained on an example. Consider the following δ -representation

$$\Delta = \langle B(d), 0, F(c), 1, G(a, b), 3, Time, 1, F(a), 2, F(d) \rangle$$

with $D_{max} = 3$ and the action

$$Time@T, G(x, y)@T_1, F(x)@T_2 \rightarrow \exists z. Time@T, G(y, z)@(T + 1), F(y)@T$$

which is applicable to Δ with the substitution $\sigma(x) = a, \sigma(y) = b$. We remove those facts from the pre-condition that do not appear in the post-condition, namely $G(a, b)$ and $F(a)$, and get an expression

$$B(d), 0, F(c), 1, -, 3, Time, 1, -, 2, F(d).$$

Next we insert the facts that appear in the post-condition and not in the pre-condition. In our case above that is the fact $G(b, n)$, where n is a fresh value. The placement of these facts is determined by the timestamps appearing in the action, which are of the form $T + d$, where T is the global time. In our example the fact $G(b, n)$ comes with the timestamp $(T + 1)$ and we get:

$$\langle B(d), 0, F(c), \infty, Time, 1, G(b, n), 2, F(d) \rangle.$$

652 after updating the truncated time differences. Notice that, for example, the relative time
653 difference between facts $F(d)$ and $Time$ is still 3.

654 However, in order to prove that actions preserve the equivalence among configurations, we
655 need yet another assumption to be able to faithfully handle time advances. The problem lies
656 within the *future facts*, that is, the facts with timestamps greater than the global time. If there
657 is a future fact P such that $\delta_{Time, P} = \infty$, then it is not the case that equivalence is preserved
658 when we advance time. For example, consider the following two configurations equivalent
659 with the upper bound $D_{max} = 3$:

$$\mathcal{S}_1 = \{Time@0, P@5\} \quad \text{and} \quad \mathcal{S}_2 = \{Time@0, P@4\}.$$

660 If we advance time on both configurations, then the resulting configurations, \mathcal{S}'_1 and \mathcal{S}'_2 , are
661 not equivalent. In particular, the truncated time difference $\delta_{Time, P}$ is still ∞ in \mathcal{S}'_1 , while it
662 changes to 3 in \mathcal{S}'_2 . Notice that the same problem does not occur neither with present nor past
663 facts, *i.e.*, the facts with timestamps that are smaller or equal to the global time.

DEFINITION 5.4

664 Given an upper bound D_{max} in a planning problem (as per Definition 5.1), a configuration \mathcal{S}
665 is called *future bounded* if for any future fact P in \mathcal{S} , the time difference $\delta_{Time, P} \leq D_{max}$.

666 Recall from Section 2 that there are two types of actions, namely, the action that advances
667 time and instantaneous actions belonging to agents. Moreover, recall that the latter actions are
668 restricted in such a way that all created facts have timestamps of the form $T + d$, where T is
669 the global time. This restriction allows us to show that actions preserve the future boundedness
670 of configurations as states the following result.

LEMMA 5.5

671 Let \mathcal{T} be a *TLSTS*, D_{max} be the upper bound in a planning problem (as per Definition 5.1),
672 and \mathcal{S} be a future bounded configuration. Let \mathcal{S}' be the configuration obtained from \mathcal{S} by
673 applying an arbitrary action in \mathcal{T} . Then \mathcal{S}' is also future bounded.

22 A Rewriting Framework and Logic for Activities Subject to Regulations

674 PROOF. Let $S \xrightarrow{a} S'$, and assume S' is not future bounded. Then there is a fact $Q'@T'$ in
 675 S' such that $T' - T > D_{max}$, where T is the timestamp of $Time$, *i.e.* the global time in
 676 both S and S' . Since S is future bounded, the fact $Q'@T'$ does not appear in S , but is created
 677 by the action a . Hence, $T' = T + D$ for some number $D \leq D_{max}$, which contradicts with
 678 $T' - T > D_{max}$. ■

679 As per Definition 5.1 the initial configuration in a planning problem is future bounded,
 680 which as per above lemma implies that all configurations in a plan are also future bounded.
 681 Notice that even if we relax the assumption that the initial configuration is future bounded,
 682 we can make it future bounded by setting the value of D_{max} to be the greater than all the
 683 timestamps in the initial configuration, *i.e.*, D_{max} would still be the upper bound on the values
 684 of the given $TLSTS$ and in the initial, goal, and critical configurations. The important result,
 685 given by the above lemma, is that future boundedness is preserved with action application.

686 Following Lemma 5.3 and Lemma 5.5, given a planning problem, we say that a δ -representa-
 687 tion is an *initial / goal / critical / future bounded δ -representation* if it is the δ -representation
 688 of an initial / goal / critical / future bounded configuration. A *plan over δ -representations* is
 689 *compliant* for a given planning problem if it does not contain any critical δ -representations
 690 and if all of its branches lead from the initial δ -representation to a goal δ -representation.

691 We are now ready to show the main result of this section.

THEOREM 5.6

692 For any given planning problem the equivalence relation between configurations given by Def-
 693 inition 5.1 is well-defined with respect to the actions of the system (including time advances)
 694 and goal and critical configurations. Any plan starting from the given initial configuration can
 695 be conceived as a plan over δ -representations.

696 PROOF. We first prove that the equivalence among configurations is well defined with respect
 697 to application of actions, *i.e.* that action application on δ -representations is unambiguous. It
 698 must be independent of the choice of configurations in (5.1). Consider the diagram below,
 699 where \mathcal{S}_1 and \mathcal{S}_2 are two equivalent configurations. Assume that \mathcal{S}_1 is transformed to \mathcal{S}'_1
 700 by means of an action α . By Lemma 5.3 the configuration \mathcal{S}_2 also complies with the time
 701 constraints required in α , and hence the action α is applicable to \mathcal{S}_2 and will transform \mathcal{S}_2 into
 702 some \mathcal{S}'_2 . It remains to show that \mathcal{S}'_1 is equivalent to \mathcal{S}'_2 .

$$\begin{array}{ccc} \mathcal{S}_1 & \xrightarrow{\alpha} & \mathcal{S}'_1 \\ \wr & & \\ \mathcal{S}_2 & \xrightarrow{\alpha} & \mathcal{S}'_2 \end{array}$$

703 We consider our two types of actions, namely, time advances and instantaneous actions (see
 704 Section 2). Let the time advance transform \mathcal{S}_1 into \mathcal{S}'_1 , and \mathcal{S}_2 to \mathcal{S}'_2 . Since only the timestamp
 705 T denoting the global time in $Time@T$ is increased by 1, and the rest of the configuration
 706 remains unchanged, only truncated time differences involving $Time$ change in the resulting
 707 configurations. Because of the equivalence $\mathcal{S}_1 \sim \mathcal{S}_2$, for a fact $P@T_1^P$ in \mathcal{S}_1 with $T_1^P \leq T$,
 708 $Time@T$ and $\delta_{P,Time} = t$, we have $P@T_2^P$ with $T_2^P \leq \hat{T}$, $Time@T$ and $\delta_{P,Time} = t$ in \mathcal{S}_2
 709 as well. Therefore, we have $\delta_{P,Time} = [t + 1]$ both in \mathcal{S}'_1 and \mathcal{S}'_2 . On the other hand for any
 710 future fact $Q@T^Q$ with $\delta_{Time,Q} = t$ in \mathcal{S}_1 and in \mathcal{S}_2 , we get $\delta_{Time,Q} = t - 1$ in both \mathcal{S}'_1 and
 711 \mathcal{S}'_2 . Therefore, \mathcal{S}'_1 and \mathcal{S}'_2 are equivalent. From Lemma 5.5, we have that both \mathcal{S}'_1 and \mathcal{S}'_2 are
 712 future bounded.

713 For the second type of actions, namely the instantaneous actions belonging to agents,
714 the reasoning is similar. Each created fact in the configuration \mathcal{S}'_1 and \mathcal{S}'_2 will be of the
715 form $P@(\mathcal{T}^1 + d)$ and $P@(\mathcal{T}^2 + d)$, where \mathcal{T}^1 and \mathcal{T}^2 represent global time in \mathcal{S}_1 and
716 \mathcal{S}_2 , respectively. Therefore each created fact has the same difference d to the global time in
717 the corresponding configuration. This implies that the created facts have the same truncated
718 time differences to the remaining facts. Hence \mathcal{S}'_1 and \mathcal{S}'_2 are equivalent. Therefore, action
719 application on δ -representations shown in (5.1) is well defined.

720 Finally, as per Lemma 5.3, \mathcal{S}_1 is a goal (respectively, critical) configuration if and only if
721 \mathcal{S}_2 is a goal (respectively, critical) configuration.

722 By induction on the length of the plan, it immediately follows that, given a planning
723 problem, any compliant plan over configurations can be represented by a compliant plan over
724 δ -representations. That is, the abstraction of configurations to δ -representations is complete.

It remains to show that the abstraction is also sound, namely that, from a compliant plan
over δ -representations for a given planning problem, we can extract a concrete plan over
configurations and that such a plan is compliant with respect to that planning problem. Any
given δ -representation corresponds to an infinite number of configurations. For example, for
the δ -representation $\langle Q_1, \delta_1, Q_2, \dots, Q_{m-1}, \delta_{m-1}, Q_m \rangle$, one of the corresponding configura-
tions is

$$\{Q_1@0, Q_2@\tilde{\delta}_1, Q_3@(\tilde{\delta}_1 + \tilde{\delta}_2), \dots, Q_m@(\tilde{\delta}_1 + \dots + \tilde{\delta}_{m-1})\}$$

725 where $\tilde{\delta}_i = \delta_i$ if $\delta_i \leq D_{max}$, and $\tilde{\delta}_i = D_{max} + 1$ if $\delta_i = \infty$. We are, however, already given
726 the initial configuration W_0 in the planning problem, for which we have $\Delta_0 = \delta_{W_0}$.

727 We prove the existence of a plan over configurations by induction on the length of the plan
728 over δ -representations. Let $\Delta_0 \rightarrow_{a_1} \Delta_1 \rightarrow_{a_2} \dots \rightarrow_{a_n} \Delta_n$ be a plan over δ -representations,
729 compliant with the respect to the given planning problem. Then Δ_0 is the δ -representation
730 of the initial configuration, *i.e.* $\Delta_0 = \delta_{W_0}$. For each $\Delta_{i-1} \rightarrow_{a_i} \Delta_i$, as per Lemma 5.3, since
731 $\Delta_{i-1} = \delta_{W_{i-1}}$, the same action a_i is applicable to the configuration W_{i-1} , resulting in W_i .
732 As proven above, and shown in (5.1), it follows that $\Delta_i = \delta_{W_i}$:

$$\begin{array}{ccccccc} \delta_{W_0} & & \delta_{W_{i-1}} & & \delta_{W_i} & & \delta_{W_n} \\ \parallel & & \parallel & & \parallel & & \parallel \\ \Delta_0 & \rightarrow_{a_1} \dots \rightarrow_{a_{i-1}} & \Delta_{i-1} & \rightarrow_{a_i} & \Delta_i & \rightarrow_{a_{i+1}} \dots \rightarrow_{a_n} & \Delta_n \\ \wr & & \wr & & \wr & & \wr \\ W_0 & \rightarrow_{a_1} \dots \rightarrow_{a_{i-1}} & W_{i-1} & \rightarrow_{a_i} & W_i & \rightarrow_{a_{i+1}} \dots \rightarrow_{a_n} & W_n \end{array}$$

733 Hence, we get a plan over configurations consisting of same sequence of actions as the
734 given plan over δ -representations. Since none of the δ -representations $\Delta_i = \delta_{W_i}$ is critical,
735 it is also the case that none of the configurations W_i is critical. Also, since $\Delta_n = \delta_{W_n}$
736 is the goal δ -representation, it follows that W_n is a goal configuration. Hence the plan
737 $W_0 \rightarrow_{a_1} W_1 \rightarrow \dots \rightarrow_{a_n} W_n$ is compliant with respect to the given planning problem. ■

738 The above theorem establishes that using δ -representations for writing plans is well defined,
739 but it does not establish a bound on the number of δ -representations. To achieve this, we
740 need the further assumption that all actions are balanced. Recall that balanced actions are
741 actions that have the same number of facts in their pre- and post-conditions. By using balanced
742 actions, the number of facts in any configuration of a plan is the same as the number of facts
743 in the plan's initial configuration. Hence, as we describe in Section 6, we can establish that
744 there is a finite number of δ -representations.

TABLE 1: Summary of the complexity results for the planning problems for balanced systems. We mark the new results appearing here with a \star .

Planning Problems	LSTSeS (No time, no branching)		TLSTSeS (Possible nonces)	
	No fresh values	Possible nonces	No branching	Possible branching
(Weak) Plan	PSPACE-complete [25]	PSPACE-complete[20]	PSPACE-complete \star	EXPTIME-complete \star
System	PSPACE-complete [25]	PSPACE-complete[20]	PSPACE-complete \star	EXPTIME-complete \star

745 6 Complexity Results

746 This section enters into the details of the complexity of the planning problems for *TLSTSeS*.
 747 These problems were introduced in [25, 21] in the setting without explicit time or branching.
 748 At the end of Section 2 we have restated these problems in our setting with explicit time and
 749 branching.

750 Recall that facts are timestamped and that there is a finite, possibly empty set of time
 751 constraints attached to a timed initial, goal and critical configuration. Recall as well that for a
 752 given initial configuration W and a finite set of goal and critical configurations, we consider a
 753 branching plan \mathcal{P} compliant if it does not contain any critical configuration, and moreover if
 754 all branches of \mathcal{P} lead from configuration W to some goal configuration.

755 Throughout this section, we assume that all actions are balanced, *i.e.*, actions have the same
 756 number of facts in their pre and post-conditions, and that the size of facts is bounded.

757 Our complexity results for the planning problems for *TLSTSeS* are summarized in 1.

758 6.1 Planning Problems for *TLSTSeS* with Non-Branching Actions only

759 We first investigate the complexity of planning problems for *TLSTSeS* when actions are non-
 760 branching and balanced and when the size of facts is bounded. We show that these problems
 761 are PSPACE-complete with respect to the parameters from the given planning problem.

762 *PSPACE-hardness:* It was shown in [20] that one can faithfully encode a Turing machine
 763 with a fixed size tape using systems with balanced actions. The same idea works in our setting
 764 with time. It is easy to modify the encoding in [20]. Timestamps do not play any important
 765 role in such encoding. Also, critical configurations are not necessary used in the encoding, so
 766 we can conclude that all three planning problems and the reachability problem for *TLSTSeS*
 767 with non-branching balanced actions and facts of bounded size are PSPACE-hard.

768 *PSPACE upper bound:* It is more interesting to show that the planning problems are in
 769 PSPACE when the size of facts is bounded and actions are non-branching and balanced. In
 770 particular, we will now use all the machinery introduced in Section 5 by using δ -representations
 771 of configurations to search for compliant plans.

772 In order to determine the existence of a compliant plan, it is enough to consider plans
 773 that never reach configurations with the same δ -configuration twice. If a plan reaches a
 774 configuration whose δ -representation is the same as a previously reached configuration, there
 775 is a cycle of actions which could have been avoided. The following lemma imposes an upper
 776 bound on the number of different δ -representations in a plan, given an initial finite alphabet.

777 Such an upper bound provides us with the maximal length of a plan one needs to consider.

LEMMA 6.1

778 Given a *TLSTS* \mathcal{T} under a finite alphabet Σ , an upper bound on the size of facts, k , and an
779 upper bound, D_{max} , on the numeric values appearing in the planning problem, namely, in \mathcal{T}
780 and in the initial, goal and critical configurations, then the number of different δ -representations,
781 denoted by $L_{\mathcal{T}}(m, k, D_{max})$, with m facts (counting repetitions) is such that

$$L_{\mathcal{T}}(m, k, D_{max}) \leq (D_{max} + 2)^{(m-1)} J^m (D + 2mk)^{mk},$$

782 where J and D are, respectively, the number of predicate symbols and the number of constant
783 and function symbols in the initial alphabet Σ .

784 PROOF. Let $\langle Q_1, \delta_{Q_1, Q_2}, Q_2, \dots, Q_{m-1}, \delta_{Q_{m-1}, Q_m}, Q_m \rangle$ be a δ -representation with m facts.
785 There are m slots for predicate names and at most mk slots for constants and function symbols.
786 Constants can be either constants in the initial alphabet Σ or names for fresh values (nonces).
787 Following [20], we need to consider only $2mk$ names for fresh values (nonces). Finally, only
788 time differences up to D_{max} have to be considered together with the symbol ∞ and there are
789 $m - 1$ slots for time differences in a δ -representation. ■

790 Intuitively, our upper bound algorithm keeps track of the length of the plan it is constructing
791 and if the length of such a plan exceeds $L_{\mathcal{T}}(m, k, D_{max})$, then the same δ -representation has
792 been reached twice. This is possible in PSPACE since the number of different δ -representations
793 given above, when stored in binary, occupies only polynomial space with respect to its
794 parameters.

795 For the below results, we assume that, given a *TLSTS* \mathcal{T} , and a finite set of goal and critical
796 configurations, it is possible to check in polynomial space whether a configuration is critical,
797 whether it is a goal configuration, and whether an action is valid, *i.e.* whether it is an instance
798 of an action from \mathcal{T} that is applicable in a given configuration.

THEOREM 6.2

799 Let \mathcal{T} be a *TLSTS* with balanced non-branching actions. Then the plan compliance problem
800 is in PSPACE with respect to m , k , and $\log_2 D_{max}$, where m is the number of facts in the
801 initial configuration, k is the upper bound on the size of facts, and D_{max} is the upper bound on
802 the numeric values appearing in the model \mathcal{T} , and in the initial, goal and critical configurations.

803 PROOF. Assume given three programs, \mathcal{C} , \mathcal{G} , and \mathcal{A} , such that they return the value 1 in poly-
804 nomial space when given as input, respectively, a configuration that is critical, a configuration
805 that contains the goal configuration, and a pair of a configuration and a transition that is valid,
806 that is, an instance of an action in the *TLSTS* \mathcal{T} is applicable to the given configuration, and
807 return 0 otherwise.

808 Let m be the number of facts in the initial configuration W . Moreover, assume as inputs
809 an upper bound, k , on the size of facts, an upper bound, D_{max} , on the numeric values
810 appearing in the planning problem, that is in the given *TLSTS* \mathcal{T} , in the initial, goal and
811 critical configurations, programs \mathcal{G} , \mathcal{C} , and \mathcal{A} , as described above, and a natural number
812 $0 \leq i \leq L_{\mathcal{T}}(m, k, D_{max})$.

813 We modify the algorithm proposed in [20] in order to accommodate explicit time. The
814 algorithm must return “yes” (*i.e.* ACCEPT) whenever there is compliant plan from the initial
815 configuration W to a goal configuration, that is a configuration S such that $\mathcal{G}(S) = 1$. In
816 order to do so, we construct an algorithm that searches non-deterministically whether such a

817 configuration *is reachable*. Then we apply Savitch's Theorem to determinize this algorithm.
 818 However, instead of searching for a plan using concrete values, we rely on the equivalence
 819 described in Section 5 and use δ -representations only. Theorem 5.6 guarantees that this
 820 abstraction is sound and faithful.

821 From \mathcal{G}, \mathcal{C} , and \mathcal{A} , it is easy to construct new functions $\mathcal{G}', \mathcal{C}'$, and \mathcal{A}' that use δ -representations
 822 instead of configurations. In particular, since time constraints associated to goal and critical
 823 configurations are also relative, these can be checked by using the truncated time differences
 824 in δ -representations.

825 The algorithm begins with W_0 set to be the δ -representation of W and iterates the following
 826 sequence of operations:

- 827 1. If W_i is representing a critical configuration, *i.e.*, if $\mathcal{C}'(W_i) = 1$, then return FAIL,
 828 otherwise continue;
- 829 2. If W_i is representing a goal configuration, *i.e.*, if $\mathcal{G}'(W_i) = 1$, then return ACCEPT;
 830 otherwise continue;
- 831 3. If $i > L_T(m, k, D_{max})$, then FAIL; else continue;
- 832 4. Guess non-deterministically an action, r , from \mathcal{T} applicable to W_i , *i.e.*, $\mathcal{A}'(W_i, r) = 1$. If
 833 no such action exists, then return FAIL. Otherwise replace W_i with the δ -representation
 834 W_{i+1} resulting from applying the action r to the δ -representation W_i . This is done as
 835 expected, by updating the facts, updating the positions of facts and the corresponding
 836 truncated time differences and continue;
- 837 5. Set $i = i + 1$.

838 We now show that this algorithm runs in polynomial space. We start with the step-counter
 839 i : The greatest number reached by this counter is $L_T(m, k, D_{max})$. When stored in binary
 840 encoding, this number takes only space polynomial to the given inputs:

$$\log(L_T(m, k, D_{max})) \leq (m - 1) \log(D_{max} + 2) + m \log(J) + mk \log(D + 2mk).$$

841 Therefore, one only needs polynomial space to store the values in the step-counter.

842 We must also be careful to check that any δ -representation, W_i , can be stored in polynomial
 843 space to the given inputs. Since our system is balanced, the size of facts is bounded, and the
 844 values of the truncated time differences are bounded, hence the size of any δ -representation,
 845 $\langle Q_1, \delta_{Q_1, Q_2}, Q_2, \dots, Q_{m-1}, \delta_{Q_{m-1}, Q_m}, Q_m \rangle$, in a plan is polynomially bounded.

846 Finally, the algorithm needs to keep track of the action r guessed when moving from one
 847 configuration to another and for the scheduling of a plan. It has to store the action that has been
 848 used at the i^{th} step. Since any action can be stored by remembering two δ -representations,
 849 one can also store these actions in space polynomial to the inputs. ■

850 The reachability problem is an instance of the plan compliance problem with an empty
 851 set of critical configurations, hence the reachability problem for *TLST*Ses with balanced
 852 non-branching actions is in PSPACE as well.

853 Next we turn to system compliance problem. Recall that besides the existence of a compliant
 854 plan it is additionally requested that no critical configuration is reachable by any sequence of
 855 actions in the given system.

THEOREM 6.3

856 Let \mathcal{T} be a *TLST*S with balanced non-branching actions. Then the system compliance
 857 problem is in PSPACE with respect to m, k , and $\log_2 D_{max}$, where m is the number of facts

858 in the initial configuration, k is the upper bound on the size of facts, and D_{max} is the upper
 859 bound on the numeric values appearing in the model \mathcal{T} , and in the initial, goal and critical
 860 configurations.

861 PROOF. In order to show that the system compliance problem is in PSPACE we modify the
 862 algorithm proposed in [25] to accommodate timestamps and time constraints. Again we rely
 863 on the fact that NPSpace, PSPACE, and co-PSPACE are all the same complexity class. We
 864 use the same notation from the proof of Theorem 6.2 and make the same assumptions. In
 865 particular, we use the algorithms \mathcal{G}' , \mathcal{C}' , and \mathcal{A}' that run in polynomial space and that check
 866 whether a timed configuration is a goal configuration, a critical configuration, or if an action
 867 is valid in the given TLSTS \mathcal{T} . Again we rely on the equivalence between configurations
 868 described in Section 5 and use δ -representations only. Theorem 5.6 guarantees us that this
 869 abstraction is sound and faithful.

870 We first need to check that none of the critical configurations is reachable from the initial
 871 configuration W . To do this we provide a non-deterministic algorithm which returns “yes”
 872 exactly when a critical configuration is reachable. The algorithm starts with W_0 set to be the
 873 δ -representation of W . For any $i \geq 0$, we first check if $\mathcal{C}'(W_i) = 1$. If this is the case, then
 874 the algorithm outputs “yes”. Otherwise, we guess an action r such that $\mathcal{A}'(r) = 1$ and that it
 875 is applicable to the δ -representation W_i . If no such action exists, then the algorithm outputs
 876 “no”. Otherwise, we replace W_i with the δ -representation W_{i+1} resulting from applying the
 877 action r to δ -representation W_i . This is done as expected, by updating the positions of facts
 878 and the corresponding truncated time differences. Following Lemma 6.1 we know that at most
 879 $L_T(m, k)$ guesses are required, and therefore we use a global step-counter to keep track of the
 880 number of actions. As shown in the proof of Theorem 6.2, the value of this counter can be
 881 stored in PSPACE.

882 Next we apply Savitch’s Theorem to determinize the algorithm. Then we swap the accept
 883 and fail conditions to get a deterministic algorithm which accepts exactly when all critical
 884 configurations are unreachable.

885 Finally, we have to check for the existence of a compliant plan. For that we apply the same
 886 algorithm as for the timed plan compliance problem from Theorem 6.2, skipping the checking
 887 of critical states since we have already checked that no critical configurations is reachable
 888 from W . From what has been shown above we conclude that the algorithm runs in polynomial
 889 space. Therefore the system compliance problem is in PSPACE. ■

890 6.2 Planning Problems for TLSTSes with possibly Branching Actions

891 We now consider the plan compliance problem when actions may be branching. In particular,
 892 we show that when actions are balanced then the plan compliance problem is EXPTIME-
 893 complete with respect to the number of facts, m , in the initial configuration, the upper bound,
 894 k , on the size of facts, the upper bound, D_{max} , on the numbers explicitly appearing in the
 895 planning problem, and the upper bound, p , on the number of post-conditions of an action. For
 896 these complexity results we use alternating Turing machines [8].

897 An alternating Turing machine (ATM) is a non-deterministic Turing machine with states
 898 that are either existential or universal states. An alternating Turing machine in an existential
 899 state accepts if *some* transition from that state leads to an accepting state, while an alternating
 900 Turing machine in a universal state accepts if *every* transition from that state leads to an
 901 accepting state. Configurations of ATMs, as with standard Turing machines, consist of a
 902 tape contents, head position and a state. Computations of alternating Turing machines can be

903 represented as trees, which is similar to the representation of branching plans in *TLSTSes*.

904 *EXPTIME-hardness*: The lower bound for the plan compliance problem can be inferred
 905 from a similar lower bound described in [26]. It was shown that one can encode alternating
 906 Turing machines by using propositional actions that are balanced and branching. Time does
 907 not play an important role for that encoding.

908 *EXPTIME upper bound*: Our upper bound algorithm uses an alternating Turing machine. In
 909 particular, we show that the plan compliance problem is in alternating-PSPACE (APSPACE)
 910 with respect to the number of facts, m , in the initial configuration, the upper bound on the
 911 size of facts, k , the upper bound, D_{max} , on the numbers appearing explicitly in the planning
 912 problem, and the upper bound, p , on the number of post-conditions of any action. That is, an
 913 alternating Turing machine can solve the plan compliance problem using polynomial space.
 914 From the equivalence between APSPACE and EXPTIME shown in [8], we can infer that the
 915 plan compliance problem is in EXPTIME with respect to the same parameters.

916 We also assume here that, given a *TLSTST*, and a finite set of goal and critical configura-
 917 tions, it is possible to check in APSPACE whether a δ -representation is a goal δ -representation
 918 or a critical δ -representation and whether an action is valid, *i.e.* whether it is an instance of an
 919 action from \mathcal{T} that is applicable in the given δ -representation.

THEOREM 6.4

920 Let \mathcal{T} be a *TLSTS* with balanced actions. Then the plan compliance problem is in EXPTIME
 921 with respect to m , k , and $\log_2 D_{max}$, and p , where m is the number of facts in the initial
 922 configuration, k is the upper bound on the size of facts, D_{max} is the upper bound on the
 923 numeric values appearing in the model \mathcal{T} , and in the initial, goal and critical configurations,
 924 and p is the upper bound on the number of post-conditions of actions in \mathcal{T} .

925 **PROOF.** We exploit the fact that the complexity classes APSPACE and EXPTIME are equiva-
 926 lent [8] and show that the plan compliance problem can be solved by an alternating Turing
 927 machine in polynomial space.

928 As with the proof of Theorem 6.2, we rely on the equivalence relation described in Section 5
 929 by using the δ -representations of configurations. Theorem 5.6 ensures that such an abstraction
 930 is sound and complete.

931 We define the following function $\text{FIND}(i, X)$, which takes a natural number, i , specifying
 932 the depth of a plan and a δ -representation, X , and returns ACCEPT if a compliant plan of
 933 depth i starting from X exists, and returns FAIL otherwise. Recall from Lemma 6.1 that it
 934 suffices to consider plans of depth bounded by $L_T(m, k, D_{max})$. Our upper bound algorithm
 935 is the following: Initialize $i = L_T(m, k, D_{max})$ and W_i as the δ -representation of the initial
 936 configuration W . Then proceed as follows:

- 937 1. If W_i is a critical δ -representation then FAIL, else continue;
- 938 2. If W_i is a goal δ -representation, then ACCEPT, else continue;
- 939 3. If $i = 0$ then FAIL, else continue;
- 940 4. Guess non-deterministically an action $X \mid \mathcal{Y} \longrightarrow_A \exists x_1.X_1 \oplus \dots \exists x_n.X_n$, that is applica-
 941 ble to W_i , yielding δ -representations $W_{i-1}^1, \dots, W_{i-1}^n$;
 942 If no such action exists return FAIL;
- 943 5. If all executions of $\text{FIND}(i-1, W_{i-1}^1), \dots, \text{FIND}(i-1, W_{i-1}^n)$ return ACCEPT, then
 944 return ACCEPT, otherwise return FAIL;

945 The fifth step is where we need the extra capabilities of an alternating Turing machine as we
 946 require that *all* executions of FIND return ACCEPT. Given the proof of Theorem 6.2 and the

947 bound, p , on the number of post-conditions of actions, it is easy to check that the alternating
948 Turing machine runs in polynomial space. ■

THEOREM 6.5

949 Let \mathcal{T} be a *TLSTS* with balanced actions. Then the system compliance problem is in
950 EXPTIME with respect to m , k , and $\log_2 D_{max}$, and p , where m is the number of facts in the
951 initial configuration, k is the upper bound on the size of facts, D_{max} is the upper bound on the
952 numeric values appearing in the model \mathcal{T} , and in the initial, goal and critical configurations,
953 and p is the upper bound on the number of post-conditions of actions in \mathcal{T} .

954 **PROOF.** Similar to proof of Theorem 6.3 we first check that a critical δ -representation is not
955 contained in any tree of actions of the system with the root W . As per Lemma 6.1 it is enough
956 to consider trees of depth bounded by $L_{\mathcal{T}}(m, k, D_{max})$.

957 For that search we define the function $\text{CHECK}(i, X)$, which takes a natural number, i ,
958 specifying the depth of a tree and a δ -representation, X and returns **ACCEPT** if a critical
959 δ -representation cannot be reached from X in a tree of depth i , and returns **FAIL** otherwise.
960 The function $\text{CHECK}(i, W_i)$ is defined as follows: We initialize $i = L_{\mathcal{T}}(m, k, D_{max})$ and set
961 W_i to be the δ -representation of the initial configuration W and proceed as follows:

- 962 1. If W_i is a critical δ -representation, then **FAIL**, else continue;
- 963 2. If $i = 0$ then **ACCEPT**, else continue;
- 964 3. Guess non-deterministically an action $X \mid \mathcal{Y} \rightarrow_A \exists x_1.X_1 \oplus \dots \oplus \exists x_n.X_n$, that is applica-
965 ble to W_i , yielding δ -representations $W_{i-1}^1, \dots, W_{i-1}^n$;
966 If no such action exists return **ACCEPT**;
- 967 4. If all of the executions of $\text{CHECK}(i-1, W_{i-1}^1), \dots, \text{CHECK}(i-1, W_{i-1}^n)$ return **ACCEPT**,
968 then return **ACCEPT**, otherwise **FAIL**.

969 The forth step is where we use the extra capabilities of an alternating Turing machine as we
970 require that *all* executions return **ACCEPT**. Consequently, it will return **FAIL** if *any* execution
971 of **CHECK** returns **FAIL**, *i.e.* if any branch reaches a critical δ -representation.

972 Then, if the function **CHECK** returned **FAIL** our upper bound algorithm stops and returns
973 **FAIL**. Otherwise the algorithm proceeds by checking for the existence of a compliant plan
974 as per algorithm given in the proof of Theorem 6.4. In case $\text{FIND}(0, W_0) = \text{ACCEPT}$ the
975 algorithm returns **ACCEPT**, and returns **FAIL** otherwise.

976 Given the proof of Theorem 6.2 and the bound, p , on the number of post-conditions of
977 actions, it is easy to check that the alternating Turing machine runs in polynomial space. Since
978 the above algorithm is in **APSPACE**, it is in **EXPTIME**. We can conclude that the system
979 compliance problem for systems with possibly branching actions is in **EXPTIME**. ■

980 As mentioned in Section 2, in addition to checking for the existence of a plan in the given
981 planning problem, we are also able to schedule a plan in all of the above cases. We take the
982 additional input j and, in the case a compliant plan exists, we output the j -th action of the
983 plan. For our **PSPACE** results from Section 6.1, we store the action for which the counter i is
984 equal to j . Since an action can be stored as two δ -configurations, we can remember the j -th
985 action in polynomial space with respect to inputs. For our **EXPTIME** results from Section 6.2,
986 we assume given the tree traversal procedure and in case the compliant plan exists, following
987 our algorithm we run the fixed traversal strategy and output the j -th action.

988 **7 Relaxing the restrictions on TLSTSeS**

989 In the previous section, we demonstrated that several problems, including the reachability
 990 problem, are decidable (PSPACE-complete or EXPTIME-complete) when assuming that
 991 actions have the following restrictions:

- 992 1. All actions are Balanced;
- 993 2. The timestamps of all facts created by an action are of the form $T + d$, where T is the
 994 current time and d a natural number;
- 995 3. Time constraints of an action are of the form show in Eq. 2.1, *i.e.*, $T_1 = T_2 \pm d$,
 996 $T_1 > T_2 \pm d$, or $T_1 \geq T_2 \pm d$, involving exactly two timestamps and a natural number d .

997 Besides the intuitions given in Section 2 for these restrictions, we show in this section that
 998 relaxing any one of these restrictions leads to the undecidability of the reachability problem.
 999 The undecidability of the reachability problem implies the undecidability of the planning
 1000 problems we study in Section 6.

1001 Kanovich *et al.* [21] have already shown that the reachability problem is undecidable when
 1002 unbalanced actions are allowed. Thus, we show that the two remaining conditions are indeed
 1003 necessarily for the decidability of the reachability problem. In Section 7.1, we demonstrate
 1004 that if we only relax condition 2 above, then the reachability problem is undecidable in general,
 1005 while in Section 7.2 we show that if we only relax condition 3 above then the reachability
 1006 is also undecidable in general. In order to obtain these undecidability results, we show that
 1007 the reachability problem can be reduced to the termination problem of a two counter Minsky
 1008 machine, which is known to be undecidable [30]. We briefly review Minsky machines:

1009 A Two-Counter Machine proposed by Minsky [30] is a machine that contains two registers
 1010 r_1 and r_2 , a set of states, \mathcal{S} , and a set of instructions, Ψ . A configuration of a Minsky machine
 1011 is a tuple $\langle k, i, j \rangle$, where k is the state of the machine, i is the value stored in the register r_1
 1012 and j the value stored in the register r_2 .

1013 There are only four types of instructions each of them leading from one state, k , to another
 1014 state, j or j_1 or j_2 , but with the following side effects on the value of registers:

- 1015 • **(Add r_i)** ins_k : $r_i = r_i + 1$; **goto** ins_j ;
- 1016 • **(Subtract r_i)** ins_k : $r_i = r_i - 1$; **goto** ins_j ;
- 1017 • **(0-test r_i)** ins_k : **if** $r_i = 0$ **goto** ins_{j_1} **else goto** ins_{j_2} ;
- 1018 • **(Jump)** ins_k : **goto** ins_j ;

1019 (1) An Add r_i instruction increments the register r_i ; (2) A Sub r_i instruction is applicable only
 1020 when r_i has a positive number and decrements it; (3) A 0-test r_i instruction is a branching
 1021 instruction leading to one state if r_i contains zero and to another state otherwise; finally (4) a
 1022 Jump instructions simply moves from one state to another without changing the values stored
 1023 in the registers. Minsky showed that the problem of determining whether a final state, a_0 , is
 1024 reachable from an initial state is undecidable. We assume, with loss of generality, that in the
 1025 initial state the registers r_1 and r_2 are set to zero.

1026 **7.1 Relaxing Advances of Timestamps**

1027 This section shows that by relaxing the restriction that timestamps of facts created by ac-
 1028 tions should be necessarily of the form $T + d$, where T is the current time of the enabling

1029 configuration and d a natural number. We generalize actions to be of the following form:

$$\begin{aligned} & \text{Time}@T, W, P_1(\vec{c}_1)@T_1, \dots, P_n(\vec{c}_n)@T_n \mid \mathcal{T} \rightarrow_A \\ & \exists u. \text{Time}@T, W, P'_1(\vec{c}'_1)@(T + f_1(T_1, \dots, T_n)), \dots, P'_m(\vec{c}'_m)@(T + f_m(T_1, \dots, T_n)), \end{aligned}$$

1030 where in the timestamps of the created facts a polynomial $f_i(T_1, \dots, T_n)$ is added to the
1031 global time, T . Polynomial $f_i(T_1, \dots, T_n)$ may contain timestamps T_1, \dots, T_n that appear
1032 as timestamps of facts in the precondition of the action. We say that such an action is
1033 *linearly-time-advancing* if all polynomials $f_i(T_1, \dots, T_n)$ are linear.

1034 Given the actions of the above form, we show that the reachability problem for these systems
1035 is undecidable already for systems with balanced actions that are linearly-time-advancing.
1036 This means that all compliance problems discussed in Section 2 are also undecidable for such
1037 systems.

THEOREM 7.1

1038 Given a *TLSTS* with balanced and linearly-time-advancing actions, the reachability problem
1039 is undecidable.

1040 **PROOF.** The proof is obtained by reducing the reachability problem of TLSTSeS with actions
1041 that are balanced and linearly-time-advancing to termination of Minsky machines. We encode
1042 an arbitrary Minsky machine M as follows:

1043 For each state label k , we associate a zero arity predicate St_k , called state fact, denoting the
1044 current state of the machine. Moreover, we use two zero arity predicates R_1 and R_2 to keep
1045 track of the value stored in the registers r_1 and r_2 , respectively. Our actions will enforce that
1046 at any given configuration there is exactly one state fact and exactly one occurrence of a R_1
1047 and a R_2 fact. We encode the values stored in the registers r_1 and r_2 , by using the timestamps
1048 of the state facts, and by using the facts R_1 and R_2 appearing in a configuration as follows:

1049 If $St_k@T$, $R_1@T_1$ and $R_2@T_2$ are the occurrences of the state fact and R_1, R_2 in a configu-
1050 ration, then such a configuration specifies that the machine is in state k , the value stored in the
1051 register r_1 is $(T_1 - T)$ and the value in r_2 is $T - T_2$. For instance, the following configuration
1052 $\{\text{Time}@7, St_a@1, R_1@3, R_2@5\}$ specifies M 's configuration $\langle a, 3, 5 \rangle$.

1053 Each of M 's instructions is encoded by the corresponding balanced linearly-time-advancing
1054 actions. The actions of the encoding have the following shape:

$$\begin{aligned} & \text{(Add } r_1) \text{Time}@T, St_k@T_1, R_1@T_2, R_2@T_3 \rightarrow_A \\ & \quad \text{Time}@T, St_j@T, R_1@(T + T_2 - T_1 + 1), R_2@(T + T_3 - T_1) \\ & \text{(Add } r_2) \text{Time}@T, St_k@T_1, R_1@T_2, R_2@T_3 \rightarrow_A \\ & \quad \text{Time}@T, St_j@T, R_1@(T + T_2 - T_1), R_2@(T + T_3 - T_1 + 1) \\ & \text{(0-test } r_1 \text{ if) Time}@T, St_k@T_1, R_1@T_2, R_2@T_3 \mid \{T_1 = T_2\} \rightarrow_A \\ & \quad \text{Time}@T, St_{j_1}@T, R_1@T, R_2@(T + T_3 - T_1) \\ & \text{(0-test } r_1 \text{ else) Time}@T, St_k@T_1, R_1@T_2, R_2@T_3 \mid \{T_1 < T_2\} \rightarrow_A \\ & \quad \text{Time}@T, St_{j_2}@T, R_1@(T + T_2 - T_1), R_2@(T + T_3 - T_1) \\ & \text{(0-test } r_2 \text{ if) Time}@T, St_k@T_1, R_1@T_2, R_2@T_3 \mid \{T_1 = T_3\} \rightarrow_A \\ & \quad \text{Time}@T, St_{j_1}@T, R_1@(T + T_2 - T_1), R_2@T \\ & \text{(0-test } r_2 \text{ else) Time}@T, St_k@T_1, R_1@T_2, R_2@T_3 \mid \{T_1 < T_3\} \rightarrow_A \\ & \quad \text{Time}@T, St_{j_2}@T, R_1@(T + T_2 - T_1), R_2@(T + T_3 - T_1) \\ & \text{(Jump) Time}@T, St_k@T_1, R_1@T_2, R_2@T_3 \rightarrow_A \\ & \quad \text{Time}@T, St_j@T, R_1@(T + T_2 - T_1), R_2@@(T + T_3 - T_1) \end{aligned}$$

1055 It is easy to show that each action faithfully encodes the corresponding instruction in M . For
1056 instance, consider the first action above encoding an **(Add** r_1) instruction. At the precondition
1057 the values stored in the registers r_1 and r_2 are, respectively, $T_2 - T_1$ and $T_3 - T_1$. In the

1058 post-condition, however, since the facts have to advance in time, the timestamp of the fact
 1059 St_j , denoting the next instruction, is changed to the current global time T . Therefore, the
 1060 timestamps of the facts R_1 and R_2 have to be updated to $T + T_2 - T_1 + 1$ and $T + T_2 - T_1$,
 1061 where the value in the register r_1 is increased by one. Also notice that **(0-test r_i)** instructions
 1062 are split into two actions: one for the case when the test is satisfied (**(0-test r_i if)** and the other
 1063 for the case when the test is not satisfied (**(0-test r_i else)**). The goal is to reach a configuration
 1064 that reaches the final state a_0 , which is encoded by the fact St_{a_0} .

1065 For soundness, the only problem could be with the action that advances the global time.
 1066 However, since all actions above take into account the global time and recompute the times-
 1067 tamps of R_1 and R_2 so that they correspond to the correct values stored in the respective
 1068 registers, the system is sound. For completeness, one can show that if we do not advance time,
 1069 the values of the timestamps of R_1 and R_2 correspond exactly to the values stored by the
 1070 registers r_1 and r_2 , since the timestamps of facts St_k , encoding instructions, are always zero.
 1071 Hence, the encoding in our system is complete.

1072 Finally, notice that we do not require critical configurations, we use only one agent A , and
 1073 no actions above updates values with fresh ones. ■

1074 Since the reachability problem is undecidable and in the proof we do not make use of any
 1075 critical states, all of the compliance problems mentioned in Section 2 are undecidable.

COROLLARY 7.2

1076 Given a *TLSTS* with balanced actions that are linearly-time-advancing, then the plan compli-
 1077 ance and the system compliance problems are undecidable.

1078 7.2 Relaxing Time Constraints

1079 Instead of relaxing the timestamps of the facts in the post-condition, we now relax the form of
 1080 time constraints in the guard of actions and investigate the complexity of the reachability prob-
 1081 lem for such systems. Recall that in our models, *TLSTSes*, time constraints are necessarily of
 1082 the form $T_1 \circ T_2 + d$, where $\circ \in \{>, \geq, =, <, \leq\}$ and d is a natural number. We relax this
 1083 condition by allowing actions to contain constraints of the form $T_1 \circ f(T_1, \dots, T_n)$, where f
 1084 is a linear polynomial and T_1, \dots, T_n are the timestamps appearing in the precondition of the
 1085 corresponding action. We call this type of actions *linearly-constrained actions*. We show that
 1086 the reachability problem for *TLSTSes* with balanced and linearly constrained actions is also
 1087 undecidable.

THEOREM 7.3

1088 Given a *TLSTS* with balanced and linearly-constrained actions, then the reachability problem
 1089 is undecidable.

1090 **PROOF.** As in the proof of Theorem 7.1, we reduce the reachability problem for *TLSTSes* to
 1091 the termination of an arbitrary Minsky Machine M .

1092 As in the proof of Theorem 7.1, the difference between the timestamps of R_1 (respectively,
 1093 R_2) and St_k will denote the value of the register r_1 (respectively, r_2). We also use the auxiliary
 1094 predicate Aux , and a predicate $Update_i^\gamma$ for each instruction γ and $i \in \{1, 2\}$ together with
 1095 the *Time* predicate to encode the effects of the instructions of M , such as the instruction to
 1096 add a value to a register. The initial configuration consists of six facts with three copies of
 1097 Aux .

$$\mathcal{I} = \{St_1@0, R_1@0, R_2@0, Aux@0, Aux@0, Aux@0, Time@0\}.$$

1098 where we assume w.l.o.g. that the values in both registers is zero. A goal configuration is any
 1099 configuration containing the facts $G = \{St_{a_0}@T_1, Aux@T_2, Aux@T_3, Aux@T_3\}$.

1100 Each of M 's instructions is encoded using a collection of auxiliary actions. Consider the
 1101 following instruction, γ , that adds the register r_1 :

(Add r_1) $ins_k: r_1 = r_1 + 1; \text{goto } ins_j$

1102 This instruction is encoded by the following four balanced and linearly-constrained actions:

- (Action γ 1)** $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, Aux@T_4, Aux@T_5, Aux@T_6 \rightarrow_A$
 $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_j@T, Update_1^\gamma@T, Update_2^\gamma@T$
- (Action γ 2)** $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_j@T_4, Update_1^\gamma@T_5 \mid \{T = T_4 + T_1 - T_3 + 1\} \rightarrow_A$
 $Time@T, R_1@T, R_2@T_2, St_k@T_3, St_j@T_4, Aux@T$
- (Action γ 3)** $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_j@T_4, Update_2^\gamma@T_5 \mid \{T = T_4 + T_2 - T_3\} \rightarrow_A$
 $Time@T, R_1@T_1, R_2@T, St_k@T_3, St_j@T_4, Aux@T$
- (Action γ 4)** $Time@T, St_k@T_3, St_j@T_4, Aux@T_5, Aux@T_6 \rightarrow_A$
 $Time@T, St_k@T_3, Aux@T, Aux@T, Aux@T$

1103 The first action is applicable only when the current state is k , specified by the fact $St_k@T_3$ in
 1104 the enabling configuration. It replaces the Aux facts with $St_j@T, Update_1^\gamma@T, Update_2^\gamma@T$.
 1105 The first fact encodes the new state j . Since the timestamp of St_j is T , whereas the values in
 1106 the registers are computed with respect to the timestamp of St_k , namely $T_1 - T_3$ and $T_2 - T_3$,
 1107 we need to update the timestamps of R_1 and R_2 to be relative to T . This is the purpose of
 1108 the facts $Update_1^\gamma@T, Update_2^\gamma@T$ and of the second and the third action. The second action
 1109 updates the timestamps of R_1 when the current time is exactly $T_4 + T_1 - T_3 + 1$, that is, the
 1110 previous value stored in the register r_1 plus 1 and relative to the timestamp of St_j . The third
 1111 action is similar and corresponds to updating the timestamp of R_2 . Only, after the second and
 1112 third action have been applied can the fourth action be enabled and applied, as this requires
 1113 two Aux facts in the precondition. The fourth action then simply forgets the previous state, k ,
 1114 by replacing St_k with Aux .

1115 The actions encoding other type of instructions are similar. We show below the encodings
 1116 of the instructions for Subtracting, instruction for the 0-test for the register r_1 and the JUMP
 1117 instruction. The remaining actions for register r_2 are similar.

γ is a **Subtract instruction for r_1** :

- (Action γ 1)** $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, Aux@T_4, Aux@T_5, Aux@T_6 \mid \{T_1 > T_3\} \rightarrow_A$
 $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_j@T, Update_1^\gamma@T, Update_2^\gamma@T$
- (Action γ 2)** $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_j@T_4, Update_1^\gamma@T_5 \mid \{T = T_4 + T_1 - T_3 - 1\} \rightarrow_A$
 $Time@T, R_1@T, R_2@T_2, St_k@T_3, St_j@T_4, Aux@T$
- (Action γ 3)** $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_j@T_4, Update_2^\gamma@T_5 \mid \{T = T_4 + T_2 - T_3\} \rightarrow_A$
 $Time@T, R_1@T_1, R_2@T, St_k@T_3, St_j@T_4, Aux@T$
- (Action γ 4)** $Time@T, St_k@T_3, St_j@T_4, Aux@T_5, Aux@T_6 \rightarrow_A$
 $Time@T, St_k@T_3, Aux@T, Aux@T, Aux@T$

γ is a 0-test instruction for r_1 :

(Action γ 1 if) $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, Aux@T_4, Aux@T_5, Aux@T_6 \rightarrow_A | T_1 = T_3$
 $Time@T, R_1@T, R_2@T_2, St_k@T_3, St_{j_1}@T, Aux@T_5, Update_1^\gamma@T$

(Action γ 1 else) $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, Aux@T_4, Aux@T_5, Aux@T_6 \rightarrow_A | T_1 > T_3$
 $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_{j_2}@T, Update_1^\gamma@T, Update_2^\gamma@T$

(Action γ 2) $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_{j_i}@T_4, Update_1^\gamma@T_5 | \{T = T_4 + T_1 - T_3\} \rightarrow_A$
 $Time@T, R_1@T, R_2@T_2, St_k@T_3, St_{j_i}@T_4, Aux@T$

(Action γ 3) $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_{j_i}@T_4, Update_2^\gamma@T_5 | \{T = T_4 + T_2 - T_3\} \rightarrow_A$
 $Time@T, R_1@T_1, R_2@T, St_k@T_3, St_{j_i}@T_4, Aux@T$

(Action γ 4) $Time@T, St_k@T_3, St_{j_i}@T_4, Aux@T_5, Aux@T_6 \rightarrow_A$
 $Time@T, St_k@T_3, Aux@T, Aux@T, Aux@T$

γ is a Jump instruction:

(Action γ 1) $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, Aux@T_4, Aux@T_5, Aux@T_6 \rightarrow_A$
 $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_j@T, Update_1^\gamma@T, Update_2^\gamma@T$

(Action γ 2) $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_j@T_4, Update_1^\gamma@T_5 | \{T = T_4 + T_1 - T_3\} \rightarrow_A$
 $Time@T, R_1@T, R_2@T_2, St_k@T_3, St_j@T_4, Aux@T$

(Action γ 3) $Time@T, R_1@T_1, R_2@T_2, St_k@T_3, St_j@T_4, Update_2^\gamma@T_5 | \{T = T_4 + T_2 - T_3\} \rightarrow_A$
 $Time@T, R_1@T_1, R_2@T, St_k@T_3, St_j@T_4, Aux@T$

(Action γ 4) $Time@T, St_k@T_3, St_j@T_4, Aux@T_5, Aux@T_6 \rightarrow_A$
 $Time@T, St_k@T_3, Aux@T, Aux@T, Aux@T$

1118 The auxiliary facts $Aux, Update_i^\gamma$ are necessary for the soundness and completeness of our
 1119 encoding. In particular, actions can only be applied in the following order:

Action γ 1, $n \times Clock$, Action γ i, $m \times Clock$, Action γ j, $h \times Clock$, Action γ 4

1120 where n, m, k is a number of time advances, possibly zero, and $\{i, j\} = \{1, 2\}$, that is, either
 1121 Action γ 1 is applied before Action γ 2 or vice-versa.

1122 We can prove by induction on the size of plans that for a given Minsky machine M and its
 1123 encoding \mathcal{T}_M as described above, then M reaches the final state a_0 if and only if \mathcal{T}_M reaches
 1124 a goal configuration from the initial configuration \mathcal{I} .

1125 Notice that we do not need any critical configurations.

1126 Since the termination problem is undecidable, so is the reachability problem for $TLSTSes$
 1127 with balanced and linearly-constrained actions. ■

1128 We can conclude that since the reachability problem is undecidable and in the proof we do
 1129 not make use of any critical states, all the compliance problems mentioned in Section 2 are
 1130 undecidable for systems with balanced and linearly-constrained actions.

COROLLARY 7.4

1131 Given a $TLSTS$ with balanced actions that are linearly-constrained, then the plan compliance
 1132 and the system compliance problems are undecidable.

1133 8 Scenario Implemented and Experimental Results Summary

1134 We implemented a small scenario simulating a visit of a subject in a clinical investigation. In
 1135 this scenario, a subject has to undergo three tests, namely, *vital signs*, *hematology*, and *urine*

1136 *tests* and in some cases a further *nephrology test*. The first three tests have to be performed
 1137 at the same day of the subject’s visit. While the vital signs and hematology tests have a
 1138 single outcome, where the data is collected, the results of the urine test may be classified in
 1139 three levels: normal, high, or very high (typically high above the reference values for some
 1140 substance). That is, the urine test has three outcomes according to the urine test result. If the
 1141 result is very high, the urine test must be repeated *within five days*, in order to make sure that
 1142 the first result is not an isolated result. Moreover, if the result of the second urine test is either
 1143 high or very high, then an extra nephrology test must be performed on the same day as the
 1144 second urine test. The visit is completed when all necessary tests have been carried out.

1145 As described in Section 3, tests are specified as a rewrite theory specifying an action,
 1146 while the time conditions in the scenario are specified using the equational theory for critical
 1147 configurations. In particular, the action for urine test has three outcomes, one for each possible
 1148 result of the test. We have also implemented the machinery described in Section 5. For this
 1149 example, it is enough to compute the canonical form whenever time advances.

1150 For our experiments using Maude, we considered the following two optimizations. Since
 1151 the order in which the leaves of a plan appear does not really matter, we can specify the
 1152 operator representing branching, $+$, to be commutative as well, by adding the attribute `comm`
 1153 to its definition. Since Maude implements rewriting modulo axioms, this reduces both the state
 1154 space and the number of solutions. The second optimization, on the other hand, follows the
 1155 lines described in [37] and involves avoiding interleavings of actions by merging (small-step)
 1156 actions into larger (big-step) actions. However, in order to be sound and complete, such a
 1157 merging of actions can only involve actions that are mutually independent. For instance, the
 1158 order in which one performs the vital signs, the hematology and the first urine test is not
 1159 important. Hence, instead of specifying each test as a different action, we can execute all three
 1160 tests as a single action. Moreover, since the urine test has three possible outcomes, while the
 1161 other test have only one outcome, the resulting (big-step) action will also have three possible
 1162 outcomes.

1163 Table 2 summarizes our main experimental results for the scenario described above when
 1164 using different parameters D_{max} as the upper-bound of numbers appearing anywhere in the
 1165 theory (see Section 5) as well as the two optimizations described above. We performed these
 1166 experiments on an Ubuntu machine (Kernel 2.6.32-37) with 3.7 Gb memory and 4 processors
 1167 of 2.67 GHz (Intel Core i5). We observed that using a commutative $+$ reduced in average the
 1168 number of states by a factor of 8, search time by a factor of 11, and the number of solutions
 1169 by a factor of 16. The use of big-step rules, on the other hand, did not affect the number of
 1170 solutions found, but reduced considerably the number of states, by a factor of 23, and search
 1171 time, by a factor 40. The accumulated reduction when using both optimizations was of a factor
 1172 58 on the number of states, 118 on search time, and 16 on the number of solutions.

1173 The Maude code for this scenario using all combinations of the two optimizations described
 1174 above as well as their experimental results can be found in [31].

1175 9 Related Work

1176 The specification of regulations has been topic of many recent works. In [6, 7, 27], a temporal
 1177 logic formalism for modeling collaborative systems is introduced. In this framework, one
 1178 relates the scope of privacy to the specific roles of agents in the system. For instance, a
 1179 patient’s test results, which normally should not be accessible to any agent, are accessible to
 1180 the agent that has the role of the patient’s doctor. We believe that our system can be adapted or

TABLE 2: Summary of our experimental results with different optimizations, *e.g.*, big-step rules and commutative +. An entry of the form $n/t/s$ denotes that the search space had a total of n states and it took Maude t seconds to traverse all states finding s solutions. DNF denotes that Maude did not terminate after 40 minutes.

D_{max}		0	2	4	6	8
Small Step	Non-Comm.	63k/91/6	166k/263/364	373k/603/4k	755k/1651/19k	DNF
	Commutative	43k/71/6	83k/119/56	141k/188/252	222k/340/792	332k/640/2k
Big Step	Non-Comm.	3k/3/6	14k/14/364	51k/67/4k	140k/220/19k	329k/508/66k
	Commutative	1k/1/6	3k/2/56	6k/5/252	13k/13/792	23k/26/2k

1181 extended to accommodate such roles depending on the scenario considered. In particular, it
 1182 also seems possible to specify in our framework the health insurance scenario discussed in [27].
 1183 De Young *et al.* describe in [11] the challenges of formally specifying the temporal properties
 1184 of regulations, such as HIPAA and GLPA. They extend the temporal logic introduced in [6]
 1185 with fixed point operators, which seem to be required in order to specify these regulations. A
 1186 temporal logic to specify regulations, such as the FDA Code of Federal Regulations (CFR),
 1187 as properties of traces abstractly representing the operations of an organization are given in
 1188 [13]. Notions of permissions and obligations are introduced to deal with regulatory sentences
 1189 as conditions or exceptions to others. An algorithm to check conformance of audit logs to
 1190 security and privacy policies expressed in a first-order logic with restricted quantification is
 1191 presented in [17]. In the case of incomplete logs a residual policy is returned.

1192 Temporal logics are suitable for specifying the temporal properties that need to be satisfied
 1193 by the traces of a system’s operation. Our approach starts with an executable specification of a
 1194 system using rewriting logic, combined with a mechanism to specify and check properties of
 1195 executions. Specifically, critical and goal configurations defined in the equational sublogic
 1196 allow us to express properties needed for generating plans for patient visits, and for monitoring
 1197 clinical investigations including FDA reporting regulations. Timestamps allow us to express
 1198 both temporal properties and timing constraints. Moreover, this approach allows us to use
 1199 existing rewriting tools, such as Maude [9], to implement our specifications and analyses.

1200 The Petri nets (PNs) community has investigated many related problems involving time. In
 1201 particular, the coverability problem of PNs is related to our partial goal reachability problem for
 1202 *TLSTS*es of a simple form - without branching actions, or critical states, or fresh values [21].
 1203 In [10], de Frutos Escrig *et al.* show decidability results for the coverability problem of a
 1204 type of Timed PNs with discrete time. There seem to be connections between our timestamps
 1205 of facts and their time (age) associated to tokens as well as connections between our time
 1206 constraints and their time intervals labeling the arcs in these PNs. However, the complexity of
 1207 their decision procedures is extremely high, as compared with our upper bounds. Notice that
 1208 branching actions and critical states are not considered there. Despite these connections, we
 1209 did not find any work that captures exactly the model presented in this paper.

1210 Real time systems differ from our setting since dense time domains, such as the real
 1211 numbers, are required, while in our intended applications, such as clinical investigations,
 1212 discrete numbers suffice. The models introduced in [1, 24, 35] deal with the specification of
 1213 real time systems and also explore the complexity of some problems.

1214 Kanovich *et al.* in [24] propose a linear logic based framework for specifying and model-

1215 checking real time systems. In particular, they demonstrate fragments of linear logic for which
 1216 safety problems are PSPACE-complete. Interestingly, their examples are all balanced which
 1217 is in accordance to some of our conditions. However, as discussed in [12], their model is
 1218 limited since one is not allowed to specify properties which involve different timestamps. In
 1219 our formalism, such properties can be specified using time constraints. In [35] conditions are
 1220 identified for which the problem of checking whether a system satisfies a property, specified
 1221 in linear temporal logic, is decidable. As their main application is for real time systems, they
 1222 also assume dense time domains, although discrete time domains can also be accommodated.
 1223 They identify non-trivial conditions on actions which allow one to abstract time and recover
 1224 completeness. We are currently investigating whether a simpler definition of balanced actions
 1225 and relative time constraints can provide more intuitive abstractions for systems with dense
 1226 times.

1227 Finally, there is a large body of work on Timed Automata (see [1] for a survey.) While we
 1228 extend multiset rewriting systems with discrete time, Timed Automaton extend automaton with
 1229 real-time clocks. There is no evident translation between our systems and Timed Automata,
 1230 but there seem to be clear correspondence between our planning problem and their reachability
 1231 problem. Moreover, the reachability problem in Timed Automata with discrete time is shown
 1232 to be PSPACE-complete [1] Although timed automaton seems suitable for modeling real-time
 1233 systems, such as circuits, it is not yet clear whether it is also suitable for modeling collaborative
 1234 systems with explicit time or the notion of fresh values.

1235 10 Conclusions and Future Work

1236 This paper introduced a model based on multiset rewriting that can be used for specifying
 1237 policies and systems which mention time explicitly. We have shown that the planning problems
 1238 for balanced systems not containing branching actions are PSPACE-complete and that the same
 1239 problems for balanced systems possibly containing branching actions are EXPTIME-complete.
 1240 We have also shown that the restrictions on the form of actions and time constraint taken in the
 1241 definition of our model, *TLSTS*, are necessary to obtain the decidability of the reachability
 1242 and planning problems.

1243 We also provided the semantics of *TLSTS*es as a linear logic with definitions theory. Our
 1244 adequacy result capitalized on the completeness of the focusing strategy for this logic.

1245 There are many directions which we intend to follow. In [32], we describe how an assistant
 1246 can help the participants of clinical investigations to reduce mistakes and comply with policies.
 1247 We are extending our current implementation into a small scale prototype in Maude in order to
 1248 collect more feedback from the health care community. One main challenge, however, is to
 1249 specify procedures in a modular fashion. One might need to specify intermediate languages
 1250 that are closer to the terminology and format used in the specification of CIs, but that are still
 1251 precise enough to translate them into a *TLSTS*. We hope that the work described in [14] may
 1252 help us achieve this goal.

1253 We would also like to extend our model to include dense times. This would allow us to
 1254 specify policies for which real-times are important. For instance, [3] describes how one can
 1255 reduce human errors by connecting medical devices and configuring them according to some
 1256 hospital policies.

1257 Another interesting problem to explore is checking whether a given plan, for example, a
 1258 plan embedded in a protocol, complies with regulations no matter how it is executed. Such
 1259 checks would help protocol design and review, and FDA audits as well as sponsors to monitor

1260 CIs and detect mistakes as early as possible.

1261 Finally, recently we have formalized Progressing Collaborative Systems that may create
1262 fresh values [22], inspired by security protocols and administrative and business processes.
1263 Such systems are efficient, *i.e.* the processes are always advancing and are completed in a
1264 bounded number of transactions. This is reflected in the complexity of the planning problems
1265 with progressing behavior. We are currently looking into extending the notion of progressing
1266 to systems with time.

1267 *Acknowledgments:* We thank Anupam Datta, Nikhil Dinesh, Deepak Garg, Insup Lee, John
1268 Mitchell, Grigori Mints, Oleg Sokolsky, and Martin Wirsing for helpful discussions.

1269 References

- 1270 [1] R. Alur and P. Madhusudan. Decision problems for timed automata: A survey. In *SFM*, pages 1–24, 2004.
- 1271 [2] J.-M. Andreoli. Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation*,
1272 2(3):297–347, 1992.
- 1273 [3] D. Arney, M. Pajic, J. M. Goldman, I. Lee, R. Mangharam, and O. Sokolsky. Toward patient safety in closed-loop
1274 medical device systems. In *Proceedings of the 1st ACM/IEEE International Conference on Cyber-Physical
1275 Systems*, ICCPS '10, pages 139–148, New York, NY, USA, 2010. ACM.
- 1276 [4] D. Baelde. *A linear approach to the proof-theory of least and greatest fixed points*. PhD thesis, Ecole
1277 Polytechnique, Dec. 2008.
- 1278 [5] D. Baelde and D. Miller. Least and greatest fixed points in linear logic. In N. Dershowitz and A. Voronkov,
1279 editors, *International Conference on Logic for Programming and Automated Reasoning (LPAR)*, volume 4790,
1280 pages 92–106, 2007.
- 1281 [6] A. Barth, A. Datta, J. C. Mitchell, and H. Nissenbaum. Privacy and contextual integrity: Framework and
1282 applications. In *IEEE Symposium on Security and Privacy*, pages 184–198, 2006.
- 1283 [7] A. Barth, J. C. Mitchell, A. Datta, and S. Sundaram. Privacy and utility in business processes. In *CSF*, pages
1284 279–294, 2007.
- 1285 [8] A. K. Chandra, D. C. Kozen, and L. J. Stockmeyer. Alternation. *J. ACM*, 28:114–133, January 1981.
- 1286 [9] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and C. Talcott. *All About Maude: A
1287 High-Performance Logical Framework*. Springer, 2007.
- 1288 [10] D. de Frutos Escrig, V. V. Ruiz, and O. M. Alonso. Decidability of properties of timed-arc petri nets. In *In
1289 ICATPN00*, pages 187–206. Springer-Verlag, 2000.
- 1290 [11] H. DeYoung, D. Garg, L. Jia, D. K. Kaynar, and A. Datta. Experiences in the logical specification of the HIPAA
1291 and GLBA privacy laws. In *WPES*, pages 73–82, 2010.
- 1292 [12] H. DeYoung, D. Garg, and F. Pfenning. An authorization logic with explicit time. In *CSF*, pages 133–145, 2008.
- 1293 [13] N. Dinesh, A. K. Joshi, I. Lee, and O. Sokolsky. Reasoning about conditions and exceptions to laws in regulatory
1294 conformance checking. In *DEON*, pages 110–124, 2008.
- 1295 [14] N. Dinesh, A. K. Joshi, I. Lee, and O. Sokolsky. Permission to speak: A logic for access control and conformance.
1296 *J. Log. Algebr. Program.*, pages 50–74, 2011.
- 1297 [15] N. A. Durgin, P. Lincoln, J. C. Mitchell, and A. Scedrov. Multiset rewriting and the complexity of bounded
1298 security protocols. *Journal of Computer Security*, 12(2):247–311, 2004.
- 1299 [16] FDA. Code of federal regulations, Title 21, Chapter 1, Subchapter D, Part 312: Investigational new drug
1300 application. Available at [http://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfcfr/
1301 CFRSearch.cfm?CFRPart=312](http://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfcfr/CFRSearch.cfm?CFRPart=312).
- 1302 [17] D. Garg, L. Jia, and A. Datta. Policy auditing over incomplete logs: Theory, implementation and applications.
1303 In *CCS'11*, 2011.
- 1304 [18] J. S. Hodas and D. Miller. Logic programming in a fragment of intuitionistic linear logic. *Inf. Comput.*,
1305 110(2):327–365, 1994.
- 1306 [19] M. Kanovich, T. B. Kirigin, V. Nigam, and A. Scedrov. Bounded memory Dolev-Yao adversaries in collaborative
1307 systems. *Inf. Comput.* Accepted for Publication.
- 1308 [20] M. Kanovich, T. B. Kirigin, V. Nigam, and A. Scedrov. Bounded memory Dolev-Yao adversaries in collaborative
1309 systems. In *FAST*, 2010.

- 1310 [21] M. Kanovich, P. Rowe, and A. Scedrov. Policy compliance in collaborative systems. In *CSF '09: Proceedings*
 1311 *of the 2009 22nd IEEE Computer Security Foundations Symposium*, pages 218–233, Washington, DC, USA,
 1312 2009. IEEE Computer Society.
- 1313 [22] M. I. Kanovich, T. B. Kirigin, V. Nigam, and A. Scedrov. Bounded memory protocols and progressing
 1314 collaborative systems. In J. Crampton, S. Jajodia, and K. Mayes, editors, *ESORICS*, volume 8134 of *Lecture*
 1315 *Notes in Computer Science*, pages 309–326. Springer, 2013.
- 1316 [23] M. I. Kanovich, T. B. Kirigin, V. Nigam, A. Scedrov, C. L. Talcott, and R. Perovic. A rewriting framework
 1317 for activities subject to regulations. In A. Tiwari, editor, *RTA*, volume 15 of *LIPICs*, pages 305–322. Schloss
 1318 Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.
- 1319 [24] M. I. Kanovich, M. Okada, and A. Scedrov. Specifying real-time finite-state systems in linear logic. *Electr.*
 1320 *Notes Theor. Comput. Sci.*, 16(1), 1998.
- 1321 [25] M. I. Kanovich, P. Rowe, and A. Scedrov. Collaborative planning with confidentiality. *J. Autom. Reasoning*,
 1322 46(3-4):389–421, 2011.
- 1323 [26] M. I. Kanovich and J. Vauzeilles. The classical ai planning problems in the mirror of horn linear logic: semantics,
 1324 expressibility, complexity. *Mathematical Structures in Computer Science*, 11(6):689–716, 2001.
- 1325 [27] P. E. Lam, J. C. Mitchell, and S. Sundaram. A formalization of HIPAA for a medical messaging system. In
 1326 S. Fischer-Hübner, C. Lambrinoukakis, and G. Pernul, editors, *TrustBus*, volume 5695 of *Lecture Notes in*
 1327 *Computer Science*, pages 73–85. Springer, 2009.
- 1328 [28] R. McDowell and D. Miller. Cut-elimination for a logic with definitions and induction. *Theoretical Computer*
 1329 *Science*, 232:91–119, 2000.
- 1330 [29] J. Meseguer. Conditional Rewriting Logic as a unified model of concurrency. *Theoretical Computer Science*,
 1331 96(1):73–155, 1992.
- 1332 [30] M. Minsky. Recursive unsolvability of post’s problem of ‘tag’ and other topics in the theory of turing machines.
 1333 *Annals of Mathematics*, 1961.
- 1334 [31] V. Nigam, T. B. Kirigin, A. Scedrov, C. Talcott, M. Kanovich, and R. Perovic. Timed collaborative systems.
 1335 <http://www2.tcs.ifi.lmu.de/~vnigam/docs/TR-TLSTS/>, June 2011.
- 1336 [32] V. Nigam, T. B. Kirigin, A. Scedrov, C. Talcott, M. Kanovich, and R. Perovic. Towards an automated assistant
 1337 for clinical investigations. In *Second ACM SIGHIT International Health Informatics Symposium*, 2012.
- 1338 [33] V. Nigam and D. Miller. Algorithmic specifications in linear logic with subexponentials. pages 129–140, 2009.
- 1339 [34] V. Nigam and D. Miller. A framework for proof systems. *J. Autom. Reasoning*, 45(2):157–188, 2010.
- 1340 [35] P. C. Ölveczky and J. Meseguer. Abstraction and completeness for Real-Time Maude. *Electr. Notes Theor.*
 1341 *Comput. Sci.*, 176(4):5–27, 2007.
- 1342 [36] P. Schroeder-Heister. Rules of definitional reflection. In M. Vardi, editor, *Eighth Annual Symposium on Logic in*
 1343 *Computer Science*, pages 222–232. IEEE Computer Society Press, IEEE, June 1993.
- 1344 [37] S. F. Smith and C. L. Talcott. Specification diagrams for actor systems. *Higher-Order and Symbolic Computation*,
 1345 15(4):301–348, 2002.