

A Normative Approach for Resilient Multiagent Systems

Geeta Mahala^{1*}, Özgür Kafali², Hoa Dam¹, Aditya Ghose¹
and Munindar P. Singh³

¹School of Computing and Information Technology, University of Wollongong, Wollongong, NSW, Australia.

²School of Computing, University of Kent, Canterbury, Kent, United Kingdom.

³Department of Computer Science, North Carolina State University, Street, Raleigh, North Carolina, USA.

*Corresponding author(s). E-mail(s): gm168@uowmail.edu.au;
Contributing authors: ozgurkafali@gmail.com; hoa@uow.edu.au;
aditya@uow.edu.au; mpsingh@ncsu.edu;

Abstract

We model a multiagent system (MAS) in socio-technical terms, combining a social layer consisting of norms with a technical layer consisting of actions that the agents execute. This approach emphasizes autonomy, and makes assumptions about both the social and technical layers explicit. Autonomy means that agents may violate norms. In our approach, agents are computational entities, with each representing a different stakeholder. We express stakeholder requirements of the form that a MAS is resilient in that it can recover (sufficiently) from a failure within a (sufficiently short) duration. We present RENO, a framework that computes probabilistic and temporal guarantees on whether the underlying requirements are met or, if failed, recovered. RENO supports the refinement of the specification of a socio-technical system through methodological guidelines to meet the stated requirements. An important contribution of RENO is that it shows how the social and technical layers can be modeled jointly to enable the construction of resilient systems of autonomous agents. We demonstrate RENO using a manufacturing scenario with competing public, industrial, and environmental requirements.

Keywords: Norms, multiagent systems, resilience

047 **1 Introduction**

048

049 Models of social interaction are central to artificial intelligence (AI) and have
050 long drawn interest from the research community, especially in the field of
051 multiagent systems (MAS) [1–4]. Socio-technical systems (STS) are a power-
052 ful way of expressing social interaction among agents and provide a way of
053 governing MAS [5, 6]. In a socio-technical system (STS), autonomous agents
054 act on behalf of the stakeholders and represent their needs [7–9]. We adopt
055 a conception of an STS [10, 11] in which agents interact with each other and
056 with the underlying technical architecture. Such an STS can be represented as
057 a multiagent system (MAS) and governed via norms [12] that regulate inter-
058 actions among the agents and help guide the agents towards the achievement
059 of their stakeholders’ needs.

060 This paper falls into the broad area of socio-technical systems. It goes
061 beyond previous research by introducing the notion of resilient socio-technical
062 systems. A resilient STS is one that seeks to meet stakeholder requirements
063 and can recover from requirement failures. In this way, resilience is an essential
064 element of the trustworthiness of an STS, especially in regards to its ability
065 [13].

066 Our intuitive reading of the resilience of an STS takes the following form:
067 *Within some deadline, if an undesirable condition, or a condition that flags a
068 departure from normal operating states, comes to pass, then within k steps a
069 desirable state (or one that represents a return to normal operating conditions)
070 can be achieved with a probability higher than a specified threshold.* We identify
071 three important aspects of resilience: (1) Resilience involves recovery from an
072 undesirable state; (2) Resilience involves recovery within a deadline (or a pre-
073 specified number of steps); (3) Resilience involves recovery with a probability
074 exceeding some (sufficiently high) threshold.

075 We tackle the above intuitions by extending STS specifications to include
076 time and quantities, providing a formalization and algorithm for automatically
077 translating STS specifications to models that can be input into model-checking
078 tools such as PRISM [14], introducing probabilistic model-checking to STS
079 specifications to reason about the probability of meeting requirements, provid-
080 ing a means to evaluate trade-offs between achieving technical objectives and
081 meeting social regulations, and finally, implementing a prototype of the entire
082 framework.

083

084 **1.1 Research Objectives and Contributions**

085

086 Our *goal* is to aid the design of resilient STSs by providing a technique and
087 tooling for evaluating alternative STS specifications with respect to the stated
088 requirements. Specifically, we have the following objectives:

089 O_1 To incorporate a rich model (time and quantities) of computational norms
090 into STS specifications as a means to regulate agent behavior and guide
091 which actions an agent should take.

092

O_2 To verify at a specified level of likelihood that an STS meets its stakeholders' requirement and can recover from a requirement failure within a specified period of time. 093
094
095

O_3 To develop techniques (supported by methodological guidelines) for leveraging probabilistic model checking to explore the trade-off space involving alternative STS designs and alternative formulations (particularly relaxations) of stakeholder requirements. 096
097
098
099

Accordingly, we present RENO (short for Resilience via Norms), a probabilistic framework that evaluates how resilient an STS specification is by verifying to what extent the MAS meets its stakeholders' requirements and evaluating the speed and extent to which the MAS recovers from a requirement failure. 100
101
102
103
104

Our contributions include (i) a formal STS specification including social norms and technical actions, and associated requirements expressed with quantities and time—achieves O_1 ; (ii) a transformation algorithm that takes a given STS specification and associated requirements, and produces a PRISM (Probabilistic Symbolic Model Checker) [14] model and associated properties in Probabilistic Computation Tree Logic (PCTL) [15, 16]—provides the means to achieve O_2 ; and (iii) a formal probabilistic verification process stating how likely a given STS specification meets stakeholder requirements and recovers from requirements failures—achieves O_2 and O_3 . 105
106
107
108
109
110
111
112
113
114

1.2 Practical Usage, as Envisioned

We envision RENO being deployed in practice via the following steps. First, the designer specifies an initial version of the STS specification. Second, the stakeholder provides the requirements expressed in PCTL. The requirements may include resilience requirements (formalized in Section 3 as system properties) where we show how the specified MAS world recovers from requirement failures. Third, RENO generates a PRISM model from the given STS specification. Fourth, the designer runs the verification process. Fifth, RENO produces an output demonstrating how likely the STS specification meets the stated requirements. Sixth, the designer refines the STS specification based on the output produced by RENO. The process iterates until the requirements are satisfied. If there is a situation where the requirements are not satisfied, the stakeholders would have the option to relax the requirements suitably. Seventh, the STS designer and stakeholders can analyse the various state variables or other parameters to produce a new relaxed requirement based on the output produced by RENO. 115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138

1.3 Organization

The rest of the document is organized as follows. Section 2 defines socio-technical systems (STSs) and introduces the relevant background. Section 3 describes the computational elements of RENO. Section 4 describes the methodological guidelines. Section 5 presents our prototype implementation

139 and verification experiments. The complete replication package, which encompasses all the artifacts and experimental data, including the outcomes of
140 actions and their corresponding execution probabilities can be accessed at [17].
141 Section 6 describes the related work. Section 7 concludes the paper.
142

143

144 2 Background

145

146 2.1 Socio-Technical Systems

147

148 Our overarching objective is to develop a framework for handling resilience
149 requirements in multiagent systems (MAS) composed of agents that are
150 afforded autonomy and self-interest and where these interests might conflict
151 with those of other agents. This notion is a departure from traditional work on
152 engineering multiagent systems, where a significant body of work assumes that
153 the agents involved have the same interests and have limited autonomy [18].

154 An important challenge in such settings is managing the aggregate
155 behaviour of the participating agents and, in particular, ensuring that the
156 aggregate behaviour meets certain requirements imposed on the MAS. One
157 effective strategy for addressing this challenge involves incorporating norms.
158 Norms constrain the behavior of individual agents in group settings (e.g.,
159 societies and communities), and regulate the interactions between those indi-
160 viduals. Norms in multiagent systems specify social controls on an agent's
161 actions and can help achieve the overall objectives of the system.

162 We conceive of a multiagent system in normative terms and, in particular,
163 as a socio-technical system (STS) consisting of a social and a technical layer.

164 The technical layer of an STS is composed of software components support-
165 ing various agent actions. The social layer comprises the stakeholders and, for
166 our purposes, the agents who represent those stakeholders. Norms regulate the
167 interactions among the entities (agents and humans) in the social layer. For
168 our purposes, the stakeholders are not formally modeled and the agents (which
169 are formally modeled) capture all the relevant actions. Specifically, the norms
170 here are directed from one agent to another and state what one agent may
171 legitimately expect of another and under what conditions. That is, the agents,
172 being autonomous parties, can violate any norm that applies to them but if
173 they do so, they are identified as being in violation of the norm. Sometimes,
174 an agent must violate a norm to achieve a greater objective [19]. The techni-
175 cal layer (actions) and social layer (norms) are specified by an STS designer,
176 in accordance with stakeholders' requirements.

177 The distinction between the social and technical layers is important
178 throughout the technical development we present below. The Methodological
179 guidelines (in Section 4) are provided that support the design and implemen-
180 tation of resilient STS. In methodological guidelines, it is demonstrated that
181 updating the STS can involve changes to both the social and technical lay-
182 ers in order to meet stakeholder requirements. Changes at the technical layer
183 involve substantial changes to agent capabilities while changes in the social
184 layer involve changes to the norms guiding agent interactions. For example, if

a given STS specification does not meet the stated requirements then either the STS specification can be refined by adding or removing agents' actions at the technical layer or adding or removing norms at the social layer.	185
	186
	187

To explain what we mean by an STS, let's describe a simple use case that illustrates how the social and technical layers arise and interplay in a practical multiagent scenario.

Example of an STS with agents, norms, and software components

Consider a scenario involving the manufacture of personal protective equipment (PPE) where different stakeholders have potentially conflicting functional and sustainability requirements. Imagine there are two textile companies for manufacturing PPE units, one (*CompanyNear*) located near a population center and the other company (*CompanyFar*) located far from the population center. This STS has four agents: (1) a textile manufacturing firm *CompanyNear*; (2) another textile manufacturing firm *CompanyFar*; (3) a *Hospital* with a requirement for varying quantities of PPE units (this demand can go up if there is an outbreak of a virus such as Covid-19); (4) an environmental *Regulator* that imposes rules governing the extent of permitted pollution and which it enforces through occasional inspections.

The companies have different kinds of environmental impacts. The *regulator* imposes stricter prohibitions on *CompanyNear* relative to *CompanyFar* since *CompanyNear* is located near the city and pollution generated by it will have a potentially greater adverse impact on the health of city residents. These norms govern this STS:

- A norm governing the amount of pollution *CompanyNear* is permitted to generate.
- A similar norm governing the amount of pollution *CompanyFar* is permitted to generate (which, in general, can be higher than for *CompanyNear* since *CompanyFar* is farther from the population center).
- Norms governing the quantities of PPE that *CompanyNear* and *CompanyFar* are committed to producing for the *Hospital* (in our example, both companies commit to producing equal quantities of PPE, but these amounts could be different in general).

Both *CompanyNear* and *CompanyFar* can act to manufacture varying quantities of PPE units (we provide details on how these actions are represented in Section 3.3). The *Hospital* agent performs only one action which involves creating demand for PPE units, which is manifested via the commitments that *CompanyNear* and *CompanyFar* make to the *Hospital*. In order to mitigate the potential increase in pollution associated with increased production of PPE units. The *CompanyNear* and *CompanyFar* agents execute actions that involve decreasing pollution. These actions are manifested via the commitments that *CompanyNear* and *CompanyFar* make to the *Regulator*. The *Regulator* agent performs a variety of actions including specifying regulatory

231 limits on the permitted pollution level for each company, occasional monitoring
232 of pollution levels, and fining firms that violate pollution limits. For the
233 purposes of our example, we focus on the specification of regulatory limits on
234 pollution, and this action manifests only in the form of prohibitions that the
235 PPE manufacturers must abide by.

236 In contrast to norms, requirements are distinct as they are specific conditions
237 or constraints that must be met to achieve a particular goal or objective.
238 In our work, we evaluate whether the stated requirement will be satisfied based
239 on the norms and actions involved. Given the above STS specification, we can
240 assert a variety of requirements. An example requirement is that both *CompanyNear* and
241 *CompanyFar* should produce 1000 units of PPE per week and
242 both companies should not exceed a pollution level of 50 ppm (parts per million).
243 Another example requirement, regarding the resilience of the STS, is
244 that both *CompanyNear* and *CompanyFar* should reduce the pollution level
245 from above 100 ppm (an unacceptable level of pollution) to below 60 ppm (an
246 acceptable level of pollution) of the chemical in question.

247

248 2.2 PCTL

249

250 Our goal is to create a practical framework to govern multiagent systems that
251 meet specified resilience requirements (along with functional requirements).
252 A practical approach to specifying such requirements is to state them in
253 probabilistic rather than absolute terms. With the probability as a tunable
254 parameter set to a high value, it is possible to require that a system be resilient
255 most of the time even if not all the time.

256 There is a growing realization in the literature that for reasons similar to
257 those that make it worthwhile to consider actions with uncertain outcomes, it is
258 important to support probabilistic goals [20]. Properties such as resilience are
259 especially amenable to a probabilistic specification since they do not concern
260 the object-level behaviour of the system (even if these must meet hard, non-
261 probabilistic, requirements).

262 Our framework emphasises STS specifications to understand the probability
263 of meeting requirements and recovering from a requirement failure within a
264 specified time. PCTL is used for specifying properties of discrete-time models
265 such as discrete-time Markov chains (DTMCs) [21]. In RENO, a PRISM model
266 is generated from the given STS specification and is also a discrete-time model,
267 whose state space is discrete and transitions are discrete steps between states.
268 Therefore, we use temporal PCTL to specify the properties of the system to
269 be checked.

270 We operationalize an STS via Probabilistic Computation Tree Logic
271 (PCTL). PCTL is derived from CTL and includes a probabilistic operator \mathcal{P}
272 [15]. PCTL is equipped with temporal operators with time bounds where time
273 is discrete and one *time unit* corresponds to one transition along an execu-
274 tion path. In PCTL, each atomic proposition is a *state formula*, which involves

275

276

assertions that are true in a single state (i.e., *state properties*). *Path formulas* describe properties of paths (i.e., sequences of states). Formally, state and path formulas in PCTL are defined as follows:

State formulas $\phi ::= \text{true} \mid a \mid \neg\phi \mid \phi \wedge \phi \mid \mathcal{P}_{\sim\text{prob}}(\Psi)$

Path formulas $\Psi ::= X\phi \mid \phi U^{\leq t} \phi \mid \phi U \phi$

Where a is an atomic proposition and Ψ represents a path formula, ϕ represents a state formula and \mathcal{P} is a probabilistic operator (details below), \sim is an inequality (i.e., $\sim \in \{<, \leq, >, \geq\}$), and $\text{prob} \in [0, 1]$ is a probability bound and $t \in \mathbb{N}$.

Path formulas Ψ use the Next (X), Bounded Until $U^{\leq t}$, and Unbounded Until (U) operators. Similar to “always \square ” and “eventually \diamond ” operators in CTL, there are temporal operators in PCTL [16]. In the literature, the temporal operators \square and G have been interchangeably used for “always” and \diamond and F have been interchangeably used for “eventually”. Here, we use G and F to align with PRISM syntax.

$$\begin{aligned} [F^{\leq t} \phi]_{\sim p} &\equiv [\text{true} U^{\leq t} \phi]_{\sim p} \\ [G^{\leq t} \phi]_{\sim p} &\equiv \neg [\text{true} U^{\leq t} \neg \phi]_{\sim (1-p)} \end{aligned}$$

A state formula is used to express the property of a model. A path formula may occur only as the parameter of the probabilistic path operator $\mathcal{P}_{\sim\text{prob}}(\Psi)$. Intuitively, a state s satisfies $\mathcal{P}_{\sim\text{prob}}(\Psi)$ if the probability of taking a path from s satisfying Ψ is in the interval specified by $\sim \text{prob}$.

PRISM is a probabilistic model checker [14]. Properties of the system are represented formally in a probabilistic temporal logic (such as PCTL) and automatically verified against an input probabilistic state transition model. PRISM takes two inputs: (i) a *probabilistic model* (such as a Markov decision process (MDP)) and (ii) a *property specification*. PRISM then conducts model checking to determine which states of the system satisfy the specification. PRISM supports path properties that can be used inside the \mathcal{P} operator such as $[F^{\leq 10} q]_{\geq 0.6}$ states that with at least 60% probability q will become true within 10 time units, while $[G^{\leq 20} r]_{\geq 0.99}$ states that with at least 99% probability r will hold at least for 20 time units.

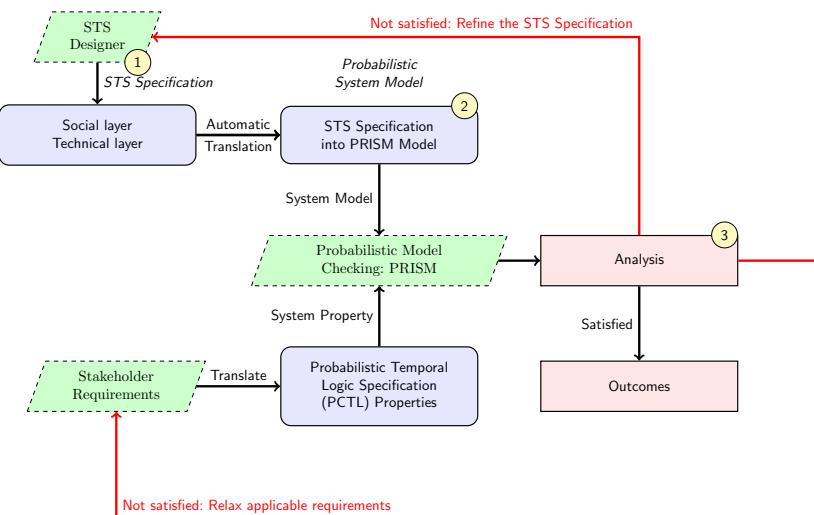
The syntax of the PRISM property specification language includes various probabilistic temporal logics, including PCTL, within PRISM. In the informal PRISM property specification language, a property is written as $\mathcal{P}_{\text{bound}}[\text{pathprop}]$, stating that the probability of being satisfied by the paths from the current state meets the bound bound .

An example of a bound would be $\mathcal{P}_{0.75}[\text{pathprop}]$ which means that the probability that the path property pathprop is satisfied by the paths from state s is greater than 0.75. To consider the nondeterministic behaviour of a system then the meaning of $\mathcal{P}_{\text{bound}}[\text{pathprop}]$ is that the probability of pathprop being satisfied by the paths from the current state meets the bound bound for all possible resolutions of the nondeterminism. Therefore, we need to reason

323 about minimum and maximum probability over all the possible resolutions of
 324 the nondeterminism. The minimum and maximum probabilities are calculated
 325 using $\mathcal{P}_{min}=?[pathprop]$ and $\mathcal{P}_{max}=?[pathprop]$, respectively. Additionally, we
 326 need to include parentheses when using logical operators in a path property
 327 (*pathprop*) because logical operators have precedence over temporal ones. An
 328 example of a property is $\mathcal{P}_{0.90}[(G \text{ ``A''}) \& (F \text{ ``B''})]$.
 329

330 3 The ReNo Framework

331 Figure 1 shows ReNO’s main components, each of which is explained below:



351 **Fig. 1** The ReNO Framework.
 352

353 **Stakeholders** Stakeholders define the requirements that the STS must satisfy.
 354 How they do so is not in our present scope but can involve a combination
 355 of creativity [22] or argumentation [23]. If it turns out that satisfying a
 356 given set of requirements is not possible, then the stakeholders would have
 357 little choice but to relax their requirements (ideally as little as possible)
 358 to make it feasible to design an STS that satisfies these (relaxed) require-
 359 ments. The reason for this, e.g., the number of steps being too small, is
 360 something that is hard to know ahead of time and is explored under the
 361 methodological guidelines as an iterative process in Section 4. In addition
 362 to the usual achievement and maintenance requirements, we introduce a
 363 new class of *resilience requirements*. Using the syntax of PCTL described
 364 above, a resilience requirement takes the following general form:
 365

$$366 \langle U \text{StateCond} \rangle \rightarrow P_{\langle ineq \rangle probabilisticConst} [F^{\langle ineq \rangle integerConst} \langle D \text{StateCond} \rangle]$$

367
 368

Here $\langle ineq \rangle$ is an operator (such as $\leq, <, \geq, or =$). $probabilisticConst$ is any real number between 0 and 1, representing a probability value.	369
Here, $integerConst$ is, as the name suggests, an integer constant, used to denote the number of steps. $UStateCond$ abbreviates a condition representing an undesirable state (such as a state in which the level of pollution exceeds an acceptable threshold). $DStateCond$ abbreviates a condition representing a desirable state (such as one in which pollution levels are below an acceptable threshold). The generic resilience requirement thus states that if an undesirable state (denoted by $UStateCond$) transpires, then the system can return to a desirable state (denoted by $DStateCond$) at some future state but before or after, depending on the nature of the inequality, the system has transitioned through $integerConst$ states, with a probability that satisfies $\langle ineq \rangle probabilisticConst$. The above equation captures the essence of resilience requirements, ensuring that the system has both the timing and the probability factors considered to facilitate the transition from an undesirable state to a desirable state. It provides a quantifiable measure for assessing the system's ability to recover and adapt in the face of undesired conditions. We provide concrete examples of resilience requirements later in the paper.	370
STS Designer The designer specifies an STS as two layers. The <i>social layer</i> describes the agents and the norms among them, whereas the <i>technical layer</i> provides the operational actions by which the agents act. The social layer consists of a set of norms that govern the interactions among the agents. Note that norms [24] in our model are directed from one party to another. In our example, each PPE manufacturer would <i>commit</i> to the <i>Regulator</i> to meet certain pollution standards. That is, the norms are pair-wise, even with multiple agents. Actions in the technical tier allow or restrict specific agent actions as they represent hard constraints. Actions describe relevant facts about the operating environment, e.g., what will (potentially) happen when an action is executed.	371
Translation of an STS Specification to a PRISM Model RENO takes the STS specification as input to generate a PRISM model. A PRISM model is a probabilistic state transition model. A probabilistic state transition model associates a <i>transition probability</i> with each transition. What we generate is a small variation in the form of <i>augmented probabilistic state transition models</i> explained in Section 3.3.1. In this case, an augmented probabilistic state transition model takes into account both the probability of selecting an action and the probability of executing an action. The process of translating an STS specification into a PRISM model is explained in Section 3.3.2.	372
Analysis RENO verifies the STS against the requirements to understand how the STS would fare, i.e., with what probability it would violate or satisfy each requirement. If the analysis suggests that refinement is required,	373
	374
	375
	376
	377
	378
	379
	380
	381
	382
	383
	384
	385
	386
	387
	388
	389
	390
	391
	392
	393
	394
	395
	396
	397
	398
	399
	400
	401
	402
	403
	404
	405
	406
	407
	408
	409
	410
	411
	412
	413
	414

415 the STS designer leverages this understanding to refine the STS specification.
 416 If the analysis determines that it is not possible to modify the STS
 417 in a way that would satisfy the requirements, the stakeholder proceeds
 418 to analyze the various state variables or parameters to identify relaxed
 419 requirements, as explained in Section 4.2.

421 3.1 ReNo Syntax

422 The ReNo uses a language for specifying norms and actions. Table 1 shows
 423 ReNo's syntax for specifying STSs. $\mathcal{AG} = \{\alpha_1, \alpha_2, \dots, \alpha_n\}$ is a finite set of
 424 agents and $\Phi = \{\text{Attr}_1, \text{Attr}_2, \dots, \text{Attr}_m\}$ is a finite set of attributes. Here,
 425 attributes are variables that are assigned values or appear in inequalities used
 426 in action specifications or in the antecedent and consequent conditions used
 427 in the various norm types (*PPE* and *pollution* in the discussion below are
 428 examples of attributes). A superscript of $+$ indicates one or more repetitions
 429 whereas superscript $*$ indicates zero or more. The operator $+=$ and $-=$ can be
 430 best understood with the following example: $x+= y$ updates x by adding y to
 431 its current value while $x-= y$ updates x by subtracting y from its current value.
 432 If y is a numeric range (NRange) as opposed to a single value, then we update
 433 x by adding (resp. subtracting) a value chosen from NRange. L_i represents the
 434 line number corresponding to the i th line in Table 1. Table 1 begins with the
 435 core concepts of ReNo and ends with primitive concepts such as Num.
 436

437 **Table 1** ReNo Syntax.

L_1	Specification	\rightarrow	$\text{Norm}^+ \mid \text{Action}^+$
L_2	Norm	\rightarrow	$\text{Commitment} \mid \text{Prohibition}$
L_3	Commitment	\rightarrow	$c(\mathcal{AG}, \mathcal{AG}, \text{Cond}, \text{Cond})$
L_4	Prohibition	\rightarrow	$p(\mathcal{AG}, \mathcal{AG}, \text{Cond}, \text{Cond})$
L_5	Action	\rightarrow	$m(\text{Cond}, \text{Stmt}^+, \text{Stmt}^+)$
L_6	Cond	\rightarrow	$\text{Expr} \mid \text{Cond} \wedge \text{Cond} \mid \text{Cond} \vee \text{Cond} \mid \sim \text{Cond}$
L_7	Expr	\rightarrow	$\top \mid \perp \mid \text{Attr} \geq \text{Num} \mid \text{Attr} \leq \text{Num} \mid \text{Attr} = \text{Val}$
L_8	Stmt	\rightarrow	$\text{Attr} += \text{NRange} \mid \text{Attr} -= \text{NRange} \mid \text{Attr} = \text{Val} \mid \text{Attr}$
L_9	NRange	\rightarrow	$\text{Num} \mid [\text{Num}, \text{Num}]$
L_{10}	Val	\rightarrow	$\text{Num} \mid \top \mid \perp$
L_{11}	Num	\rightarrow	Any numeric literal (drawn from \mathcal{R})

449
 450 L_8 Stmt can be either the update of the value of Attr using them $+=$ or $-=$
 451 operators, the assignment of a value val to Attr or a reference to Attr
 452 in isolation. For example, if Stmt is $\text{PPE}+= [100, 220]$ where PPE is an
 453 attribute, then the current value of PPE is increased by adding a value
 454 between 100 and 220.

455 L_7 An expression Expr evaluates to the logical true or false constant or an
 456 inequality involving Attr and Num.

457 L_6 A condition Cond can be an expression (Expr) or a conjunction or
 458 disjunction of expressions.

459 L_5 An action m consists of condition Cond and two lists (as discussed later,
 460 these are the DeleteList and AddList, respectively). The elements of each

list are comma-separated. An example of an action is $m(true, \{PPE, pollution\}, \{PPE+=[50, 100], pollution+=PPE*[0.2, 0.4]\})$. Here the condition is the logical constant true. The DeleteList consists of the prior values of PPE and pollution. The AddList consists of the updated values of PPE and pollution. For example, let's assume the current state s has a PPE value of 10. After executing the action, the new value of the state variable PPE can be updated with any value in the range $[50, 100]$, plus the previous value. If we consider a PPE value of 60 from the range, the next PPE value in the next state s' would be 70.	461 462 463 464 465 466 467 468 469
L_3 - L_4 A commitment consists of two agents (debtor and creditor) and a pair of conditions, with the debtor promising the creditor that it will make the second condition true if the first condition is made true. Similarly, a prohibition involves the debtor agent promising the creditor agent to not let the second condition become true if the first condition is made true. Consider the following example prohibition: $p(CompanyNear, Regulator, true, pollution \geq 60)$. Here, the debtor is CompanyNear and the creditor is Regulator. CompanyNear promises Regulator that it will always (the first condition is always true) ensure that the second condition does not become true (pollution exceeding 60ppm).	470 471 472 473 474 475 476 477 478 479
L_2 A norm may be a prohibition or a commitment.	480
L_1 An STS specification consists of one or more norms and one or more actions.	481 482
We will use Listing 1 to explain the technical and social layers of an STS specification in the following sections.	483 484 485
3.2 Social Layer: Norms	486 487
Following Kafali et al. [11] and Singh [12], we define a <i>norm</i> as a tuple $\langle n, SBJ, OBJ, ant, con \rangle$, where n represents the norm type from $\{c, p\}$; $SBJ \in \mathbb{AG}$ is its subject; $OBJ \in \mathbb{AG}$ is its object; ant and con are conditions (Cond in Table 1 above) that represent the antecedent and consequent of the norm, respectively. The set of norm types $\{c, p\}$ consists of <i>commitments</i> (c) and <i>prohibitions</i> (p).	488 489 490 491 492 493
A commitment indicates that the subject is committed to its object to making the consequent true if the antecedent holds. Consider the following commitment from Listing 1: $C_1(CompanyNear, Hospital, true, PPE \geq 100)$. The CompanyNear is committed to the Hospital to always (note that the antecedent is true) producing 100 units of PPE or more. The company is the “debtor” for this commitment, and it would violate its commitment if it fails to produce the specified amount of PPE.	494 495 496 497 498 499 500
A prohibition (p) indicates that its subject is prohibited by its object from making the consequent true when the antecedent holds. Consider the following prohibition from Listing 1: $P_1(CompanyNear, Regulator, true, pollution \geq 60)$. For example, the regulator prohibits the company from polluting above 60 ppm at any time. CompanyNear would violate its prohibition if pollution goes above the specified level of the chemical in question.	501 502 503 504 505 506

507 As seen in Listing 1, we have used a subscript number with each norm
508 type, where the subscript denotes the instance of each norm type (e.g., C_1 and
509 C_2 are two instances of commitments). For example, in Listing 1 each com-
510 pany has one commitment, so we have defined commitment $C_1(\text{CompanyNear},$
511 $\text{Hospital}, \text{true}, \text{PPE} \geq 100)$ for CompanyNear and $C_2(\text{CompanyFar}, \text{Hospital},$
512 $\text{true}, \text{PPE} \geq 100)$ for CompanyFar . Similarly, each company has one prohibi-
513 tion, so we have defined $P_1(\text{CompanyNear}, \text{Regulator}, \text{true}, \text{pollution} \geq 60)$
514 for CompanyNear and $P_2(\text{CompanyFar}, \text{Regulator}, \text{true}, \text{pollution} \geq 80)$ for
515 CompanyFar .

516

517 3.3 Technical Layer: Actions and State Transitions

518

519 The technical layer consists of a finite set of operational actions $A =$
520 $\{a_1, a_2, \dots, a_k\}$ for each agent to execute. In our use case, CompanyNear can
521 choose to manufacture 100 PPE units ($a11$) or 120 PPE units ($a12$) and sim-
522 ilarly, CompanyFar can choose to manufacture 100 PPE units ($a21$) or 120
523 PPE units ($a22$).

524 An action is represented as $m(\text{condition}, \text{DeleteList}, \text{AddList})$, where con-
525 dition is a Cond (in the sense of Table 1) while each of AddList and DeleteList
526 is a List (also in the sense of Table 1). If an action is applicable (i.e., the condi-
527 tion holds in the current state), the action is executed, leading to a transition
528 to a new state where the old values of the attributes contained in DeleteList are
529 removed and the new values of the attributes specified in AddList are added.
530 The outcome of an action can be either deterministic or nondeterministic. The
531 AddList can specify deterministic or nondeterministic outcomes. The use of a
532 range indicates that the outcome of the execution of that action can lead to a
533 state where the variable in question is assigned any value within its range. In
534 Listing 1 below, the execution of the action $a11$ can lead to a state where the
535 new value of the state variable PPE can be any value in the range $[50,100]$ and
536 the new value of the state variable pollution is the value of PPE multiplied by
537 any step-size in the real interval $[0.2, 0.4]$.

538

539 3.3.1 Selection probability

540

541 The RENO should support agent autonomy; agents are free to select which
542 action they wish to execute, but this choice is informed by the applicable
543 norms. We would like to select an action which has the most compliant
544 behaviour. We consider how likely an action is compliant to each norm. Because
545 an action may be more compliant to a particular norm but at same time, not
546 compliant to any other norms. Our framework RENO helps the agent to select
547 an action from a set of actions that has the most compliant behaviour. For
548 instance, in a given state, CompanyFar can execute either $a21$ and $a22$ to man-
549 ufacture 100 PPE. In a given state, let's assume that there is a 100% chance
550 of action $a21$ complying with a commitment, while there is a 10% chance of
551 action $a21$ complying with a prohibition. Similarly, let's assume that there is an
552 50% chance of action $a22$ complying with a commitment, while there is an

Listing 1 STS specification for PPE manufacturing.

1	$P_1(CompanyNear, Regulator, true, pollution \geq 60)$	553
2	$C_1(CompanyNear, Hospital, true, PPE \geq 100)$	554
3	$P_2(CompanyFar, Regulator, true, pollution \geq 80)$	555
4	$C_2(CompanyFar, Hospital, true, PPE \geq 100)$	556
5		557
6	$a11: m(true, \{PPE, pollution\}, \{PPE+=[50, 100], pollution+=PPE*[0.2, 0.4]\})$	558
7	$a12: m(true, \{PPE, pollution\}, \{PPE+=[80, 120], pollution+=PPE*[0.5, 0.7]\})$	559
8	$a10: m(true, \{PPE, pollution\}, \{PPE+=[0, 0], pollution+=PPE*[0.0, 0.0]\})$	560
9		561
10	$a21: m(true, \{PPE, pollution\}, \{PPE+=[50, 100], pollution+=PPE*[0.2, 0.4]\})$	562
11	$a22: m(true, \{PPE, pollution\}, \{PPE+=[80, 120], pollution+=PPE*[0.5, 0.7]\})$	563
12	$a20: m(true, \{PPE, pollution\}, \{PPE+=[0, 0], pollution+=PPE*[0.0, 0.0]\})$	564

80% chance of action $a21$ complying with a prohibition. In this scenario, the RENO will select action $a21$, which demonstrates the highest level of compliant behavior by considering both commitment and prohibition for execution in that particular state.

Therefore, we wish to model both agent choice (i.e., which action to execute) and transition probabilities (uncertain outcomes accruing from the execution of an action). We make a distinction between *selection probability* and *execution (or transition) probability*. The selection probability is important because we use *selection probability* to determine the norm-compliant behaviour.

We use the common mathematical form “softmax” to map weights to probabilities. Softmax or normalized exponential function is a mathematical function used to convert a vector of \mathcal{R} real numbers into a probability distribution of \mathcal{R} possible outcomes. The softmax function normalizes the input vector, making sure that the resulting values range between 0 and 1 and that their sum adds up to 1, which is a requirement for a probability distribution. Consequently, the output vector obtained from softmax can be interpreted as a probability distribution over the different classes or categories. The essential intuition is that selection is a means to deal with autonomy while execution probabilities are a way to deal with a stochastic environment. Here, a stochastic environment means one in which instructing an agent to take a particular action does not necessarily lead to that action being carried out.

We define a function $prob(a_k, s, A, N)$ (Equation 1) which calculates the likelihood of an agent choosing an action from a given state (i.e., selection probability). This function computes the propensity of an agent selecting an

599 action for execution, given the impact of one execution of the action on com-
 600 pliance with the applicable set of norms. When propensity is computed for all
 601 actions, we transform the resulting probabilities into a probability distribution
 602 for the set of actions.

603

604
$$prob(a_k, s, A, N) = \prod_{i=1}^{|N|} \frac{w_i^{sat(a_k, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} \quad (1)$$

605

606 where $prob(a_k, s, A, N)$ indicates the probability of selecting action a_k from
 607 a set of actions A in a state s , $sat(a_k, n_i)$ indicates the probability that action
 608 a_k satisfies the current target set out by norm n_i , and w is a parameter that
 609 indicates our willingness to permit the norms that instantiate n_i to be violated
 610 (e.g., a higher w indicates that a norm violating action is less likely to be
 611 picked). The weight w_i is not a part of the norm n_i . The willingness to violate
 612 the norm could be contingent on the agent, the current state or the potential
 613 norm-violating action, or all three. We do not propose to be prescriptive here
 614 about which intuition needs to be adopted. Indeed, any of these intuitions
 615 might be valid. We simply provide machinery in this equation to obtain a
 616 selection probability based on the willingness to violate the norm.

617

618 Now we will describe the process of computing $sat(a_k, n_i)$. The Equation 2
 619 is used to capture the likelihood of an action being compliant with prohibi-
 620 tions. Equation 3 captures the likelihood of an action being compliant with
 621 commitments.

622 For an action a_k , a norm n_i , and an attribute Attr that is common to a_k
 623 and n_i , we set the probability of a_k satisfying n_i to 0 if the upper-bound for
 624 Attr as described in a_k is less than the current target for Attr as described
 625 in n_i (as for a commitment) or the lower-bound for Attr as described in a_k
 626 is greater than or equal to the current target for Attr as described in n_i (as
 627 for a prohibition). Similarly, we set the probability of a_k satisfying n_i to 1 if
 628 the lower-bound for Attr as described in a_k is greater than or equal to the
 629 current target for Attr as described in n_i (commitment) or the upper bound
 630 for Attr as described in a_k is less than the current target for Attr as described
 631 in n_i (prohibition). Note that by computing $sat(a_k, n_i)$ as a probability, we
 632 eliminate cases where (unnecessarily) increasing or decreasing the amount of
 633 a given attribute beyond the target value might have an undesired effect on
 634 calculating the likelihood of a given action. For example, the probability of
 635 satisfaction is the same if the action produces exactly the target value or twice
 636 the target value. If the current target for Attr as described in n_i is within
 637 the interval for Attr as described in a_k , then we compute the probability of
 638 satisfaction based on a uniform distribution of values from the interval for Attr
 639 as described in a_k . For example, if 40% of the values for Attr within the lower
 640 bound and upper bound satisfies the current target, then $sat(a_k, n_i)$ is 0.4.

641

642 $c_target(\text{Attr}) = n_i_target(\text{Attr}) - current(\text{Attr})$

643

644

$$sat(a_k, n_i) = \begin{cases} 1, & \text{if } u_{Attr} \leq c_target(Attr) \\ 0, & \text{if } l_{Attr} \geq c_target(Attr) \\ \frac{c_target(Attr) - l_{Attr}}{u_{Attr} - l_{Attr}}, & \text{otherwise} \end{cases} \quad \begin{array}{l} 645 \\ 646 \\ 647 \\ 648 \end{array}$$

$$sat(a_k, n_i) = \begin{cases} 1, & \text{if } l_{Attr} \geq c_target(Attr) \\ 0, & \text{if } u_{Attr} \leq c_target(Attr) \\ \frac{u_{Attr} - c_target(Attr)}{u_{Attr} - l_{Attr}}, & \text{otherwise} \end{cases} \quad \begin{array}{l} 649 \\ 650 \\ 651 \\ 652 \\ 653 \end{array}$$

where $n_i.target(Attr)$ is the target value for the consequent of norm n_i , $current(Attr)$ is the value of the attribute in the current state, l_{Attr} and u_{Attr} are lower bound and upper bound of attribute Attr in the DeleteList or AddList of a_k . 654
655
656
657

Now, we provide a proof for the Equation 1 that it provides a probability distribution for the set of actions in a given state s . 658
659
660

Theorem 1 *To prove that the sum of probabilities of all actions for state s is equal to 1, we need to sum the probabilities of all actions over the set of available actions A :* 661
662

$$\sum_{a_k \in A} prob(a_k, s, A, N) = 1 \quad \begin{array}{l} 663 \\ 664 \\ 665 \end{array}$$

Proof Let's substitute the value of $\sum_{a_k \in A} prob(a_k, s, A, N)$ from equation 1. 666

$$\sum_{a_k \in A} prob(a_k, s, A, N) = \sum_{a_k \in A} \prod_{i=1}^{|N|} \frac{w_i^{sat(a_k, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} = \prod_{i=1}^{|N|} \sum_{a_k \in A} \frac{w_i^{sat(a_k, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} \quad \begin{array}{l} 667 \\ 668 \\ 669 \\ 670 \end{array}$$

$$\prod_{i=1}^{|N|} \sum_{a_k \in A} \frac{w_i^{sat(a_k, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} = \prod_{i=1}^{|N|} \frac{\sum_{a_k \in A} w_i^{sat(a_k, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} \quad \begin{array}{l} 671 \\ 672 \\ 673 \\ 674 \end{array}$$

Now let's take $\frac{\sum_{a_k \in A} w_i^{sat(a_k, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}}$ from equation B and using the properties of summation, we can split the sum into two parts. 675
676
677

$$\frac{\sum_{a_k \in A} w_i^{sat(a_k, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} = \frac{w_i^{sat(a_1, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} + \frac{w_i^{sat(a_2, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} + \dots + \frac{w_i^{sat(a_n, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} \quad \begin{array}{l} 678 \\ 679 \\ 680 \end{array}$$

Now, let's bring the common denominator, $\sum_{a \in A} w_i^{sat(a, n_i)}$, inside each term: 681
682

$$\frac{w_i^{sat(a_1, n_i)} + w_i^{sat(a_2, n_i)} + \dots + w_i^{sat(a_n, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} \quad \begin{array}{l} 683 \\ 684 \end{array}$$

Since the denominator $\sum_{a \in A} w_i^{sat(a, n_i)}$ is the sum of the exponential values of all actions in the A so it can be canceled out. 685
686
687

$$\frac{w_i^{sat(a_1, n_i)} + w_i^{sat(a_2, n_i)} + \dots + w_i^{sat(a_n, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} = 1 \quad \begin{array}{l} 688 \\ 689 \\ 690 \end{array}$$

691 So,

$$\frac{\sum_{a_k \in A} w_i^{sat(a_k, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} = 1 \quad (C)$$

696 Now put value of $\frac{\sum_{a_k \in A} w_i^{sat(a_k, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}}$ from equation C in equation B.

$$\prod_{i=1}^{|N|} \frac{\sum_{a_k \in A} w_i^{sat(a_k, n_i)}}{\sum_{a \in A} w_i^{sat(a, n_i)}} = \prod_{i=1}^{|N|} *1 = 1$$

700 Hence,

$$\sum_{a_k \in A} prob(a_k, s, A, N) = 1$$

703 \square

704 3.3.2 Translation of an STS specification to a PRISM model

706 The PRISM model is constructed as a probabilistic state transition model.
707 In the following algorithms, c_target(Attr), n_target(Attr), and current(Attr)
708 have the same denotations as in Equation 1, Equation 2, and Equation 3. In
709 Algorithm 1, we compute a PRISM model that includes a transition probability
710 for every feasible transition $\langle s, s' \rangle$ where s' is one of the states that could
711 result from executing action a in state s . We also note that although our state
712 schema includes a variable that denotes the agent whose turn to act, this is
713 entirely an artifact of the round-robin agent execution technique we are using
714 for approximately checking an STS specification. In the algorithms defined
715 below, we can ignore the state variable called *agent* except in situations where
716 we explicitly refer to it.

717 We also note that for many of the actions we specify their resulting states
718 in terms of the deltas (i.e., the operations performed) on the prior values of the
719 state variables. In the algorithms below, we avoid this complexity by simply
720 assuming that the actions are specified in terms of the values of the state
721 variables in the resulting state instead of explicitly describing the changes
722 or operations applied to the state variables. Thus, the algorithms avoid the
723 complexity of detailing how the state variables are updated. They focus on the
724 end result, rather than the specific steps involved in achieving that result.

725 Algorithm 1 translates an STS specification into a PRISM model. The
726 algorithm takes two inputs: 1) An STS specification consisting of a set of
727 agents, where each agent has a set of norms N and a set of actions A . 2) The
728 set of action execution probabilities given by ACTIONEXECPROB (these are
729 the probabilities associated with the action execution transitions). In practical
730 settings, we expect these probabilities to be obtained from past execution data
731 (for this paper's experimental purposes, these values are randomly generated).

732

- 733 • We create the initial state s_0 by assigning 1 to variable *agent* and assigning
734 random values to other state variables (Line 2).
- 735 • The consequences of actions always lie in that range. The number of
736 outcomes of each action is calculated based on a step size. If we use a small

Algorithm 1 Generating PRISM model from an STS specification.	737
Input: An STS specification consisting of a set of agents, where each agent is specified via:	738
A set of norms N	739
A set of actions A	740
Input: w is a user defined parameter that indicates our willingness to permit the norms that instantiate n where $n \in N$ to be violated.	741
Input: A function $\text{ACTIONEXECPROB} : A \times S \rightarrow \mathbb{R}$ which is the probability of arriving in a given state in S by executing an action in A , independent of which state we are in when this action is executed.	742
Output: A PRISM model $(S, A, \text{TransitionProb})$ where S is a set of states, A is a set of actions and a function $\text{TransitionProb} : S \times S \rightarrow \mathbb{R}$	743
1: Create an initial state s_0 by assigning to <i>agent</i> the ID of the first agent in the input STS specification and assigning the value 0 to all the other state variables (note that all of these are necessarily numeric valued).	744
2: $S := \{s_0\}$	745
3: while $S \neq \emptyset$ do	746
4: Let $s \in S$	747
5: for each $a \in A$ for the agent denoted by the current value of <i>agent</i> do	748
6: if <i>agent</i> \neq last agent in the STS specification then	749
7: s' is obtained from s by removing value assignments to variables that appear in the DeleteList of a and adding variable value pairs that appear in the AddList of a ;	750
8: <i>agent</i> := next agent in the STS specification.	751
9: else if <i>agent</i> = last agent in the STS specification then	752
10: s' is obtained from s by removing value assignments to variables that appear in the DeleteList of a and adding variable value pairs that appear in the AddList of a ;	753
11: <i>agent</i> := first agent in the STS specification.	754
12: end if	755
13: $\text{TransitionProb}(s, s') :=$	756
$\text{ACTIONSELECTIONPROBABILITY}(s, w, a, A, N)$	757
$* \text{ACTIONEXECPROB}(a, s')$	758
14: end for	759
15: $S := S - \{s\}$	760
16: end while	761

value for step-size, then we have more number of outcomes of an action as compared to a large value for step-size. For example, the consequences of the action a_{11} always lie between 50 and 100. If the STS designer chooses the step size as 10 then this action has six consequences, i.e., 50, 60, 70, 80, 90, and 100 whereas if the step size is 20 then it has only three consequences, i.e., 50, 70, and 90. It can be noticed that action a_{11} has six consequences when step-size is 10 whereas action a_{11} has three consequences when step-size is 20.

774
775
776
777
778
779
780
781
782

- In Line 5, when a specific agent is selected in a state then all transitions from this state lead to the states where we execute the set of actions (determined by the non-zero value of an agent, which serves as an identifier for the selected agent). Then, we add transitions corresponding to all non-deterministic outcomes of executions of the actions identified by the value of the *agent* variable (Line 6 or Line 9).
- We build a PRISM model as a probabilistic state transition model. Hence, we need to compute the transition probability for each state transition. Each state's state transition probability is calculated by multiplying the selection probability of an action we execute with the associated probabilities for that action as defined by the input ACTIONEXECPROB (Algorithm 1). For a state where a specific agent is selected, we calculate the action selection probability for each action in the set of actions A for that agent using Equation 1.
- Algorithm 2 is used to calculate an action selection probability in a given state.
 - We go through each action a' in the set of actions A (Line 2 in Algorithm 2).
 - Then for each action a' , we go through each norm n in the set of norms N to compute action a' 's probability which satisfies the current target set out by the norm n (Line 3 in Algorithm 2).
 - If norm n is prohibition then we use line 4 to line 7 in Algorithm 2 to compute action a' 's probability which satisfies the current target set out by the norm prohibition.
 - Similarly, if norm n is commitment then line 8 to line 11 in Algorithm 2 are used to compute the actions a' that satisfies the current target set out by the norm commitment. For example, agent CompanyNear is selected in the current state s . Then, we calculate the action selection probability for each action $a11$, $a12$, and $a10$ based on the norms P_1 and C_1 (from Listing 1).
 - Finally, we get the action selection probability for each action using Equation 1 (Line 14 in Algorithm 2).
- Finally, Line 13 in Algorithm 1 will be used to compute the transition probability for the updated state which is added using either Line 7 or Line 10 in Algorithm 1.
- For example, the current agent is CompanyNear in state s . For each action, such as $a11$, the state variables are updated (which results in the current state being updated to the next state) and state transition probability for the next state is computed in Line 13 in Algorithm 1.
- Line 13 computes the state transition probability which is equal to the multiplication of the selection probability of $a11$ with its action execution probability. This process is repeated for the remaining actions, i.e., $a12$ and $a10$. As discussed in Section 3.3, the states resulting from these transitions are obtained from the corresponding prior states by removing

the value assignments to variables in the DeleteList and adding the new values to these variables as specified in the AddList of the action.	829
	830
	831
	832
	833
	834
	835
	836
	837
1: function ACTIONSELECTIONPROBABILITY(s, a, w, A, N) where a is the action of interest (i.e., the action whose probability of selection we wish to compute), s is the current state and N is the currently applicable set of norms	838
2: for each $a' \in A$ for the applicable agent in s do	838
3: for each $n \in N$ for the applicable agent in s do	839
4: if n is a prohibition and state variable Attr is common to a and n then	840
5: $c_target(Attr) := n_target(Attr) - current(Attr)$	841
6: $sat(a', n)$ is assigned the value computed using Equation 2	842
7: end if	843
8: if n is a commitment and state variable Attr is common to a and n then	844
9: $c_target(Attr) := n_target(Attr) - current(Attr)$	845
10: $sat(a', n)$ is assigned the value computed using Equation 3	846
11: end if	847
12: end for	848
13: end for	849
14: ACTIONSELPROB(s, a, A, N) := action selection probability computed using Equation 1.	850
15: return ACTIONSELPROB(s, a, A, N)	851
16: end function	852

Note that we assume that each action and each norm refers to only one state variable (i.e., only one Attr).

Example of PRISM model generated using Algorithm 1. It would be easier to visualise the model graphically. In PRISM, the transition matrix of the model can be exported in the Dot (Dot is a graph description language) format, which allows easy graphical visualisation of the graph structure of the model. Figure 2 is constructed based on the transition matrix of a PRISM model generated using Algorithm 1 from a simplified STS specification with two agents where each agent has two actions. Figure 2 shows part of a PRISM model but the complete diagram is provided in the supplementary material. We describe the application of our algorithm from State 31 (2,3,1)—the red box in the diagram, where agent = 2, PPE = 3, and pollution = 1. Since the agent = 2, the CompanyNear agent is selected as a current agent. Two actions are executed next and each action has two consequences. In State 31, there are four potential state transitions: with 0.11 probability PPE is increased by one in the next state (State 1—green box) where agent = 1; with 0.235 probability

	857
	858
	859
	860
	861
	862
	863
	864
	865
	866
	867
	868
	869
	870
	871
	872
	873
	874

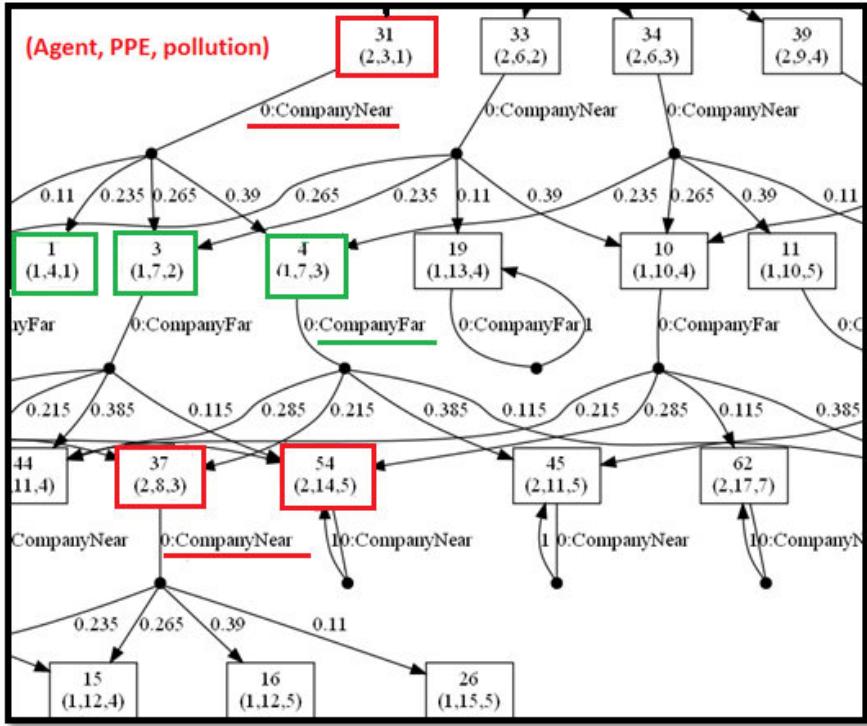


Fig. 2 Sample (partial) PRISM model using Algorithm 1.

PPE is increased by four and pollution is increased by one where agent = 1 in the next state (State 3 green box); with 0.265 probability PPE is increased by four and pollution is increased by two where agent = 1 in the next state (State 4 green box). In state 4 (1,7,3)—the green box in the diagram where agent = 1; PPE = 7; and pollution = 3. Since the agent = 1, the CompanyFar agent is selected as a current agent. There are two actions for CompanyFar to be executed next. There are four resulting states, such as 37 (2,8,3) and 54 (2,14,5). Then, state 37 (2,8,3) delivers agent = 2 so CompanyFar is again selected using the round-robin agent execution technique.

3.4 Requirements as PRISM properties

RENO supports three core types of requirements: achievement and maintenance requirements are essential to the norm literature (e.g., a commitment to achieve something to a certain level or a prohibition to maintain something at a certain level). The third type is the resilience requirements (novel to RENO) to verify that an STS can recover from requirement failures. We believe these three types of requirements cover realistic verification scenarios, to help guide STS designers. In the following, we consider these three types of requirements and their formalization in PRISM syntax.

Achievement and Maintenance. Consider the requirement that the probability that the company can manufacture at least X PPE units within K timesteps and maintain the pollution level below Y ppm should be no less than P. We formalize this in PRISM as follows:

$$P_{min=?}[(pollution < Y) \quad \mathbf{U}^{\leq K} \quad (PPE \geq X)]$$

If we achieve minimum probability P ($P_{min=?} = P$), then this requirement is satisfied.

Resilience Consider a situation where we find the pollution level above some threshold (say X ppm). Then we wish to achieve a better state (where the over-pollution problem has been resolved by bringing it down to below Y ppm) within K steps to be no worse than some threshold value.

We formalize this requirement in PRISM as follows:

$$filter(print, P_{min=?}[\mathbf{F}^{\leq K} pollution \leq Y], pollution \geq X)$$

filter() is a PRISM device that allows us to perform max or min over multiple states satisfying some property (here that property is $pollution \geq X$). If we achieve a min of 0, that tells us that we are not resilient at all while if we achieve a minimum of 1.0, then we are fully resilient.

For resilience, we are interested in ensuring that our system remains within the states that satisfy the above property with a minimum probability of 0.75. We formalize the second part of the requirement in PRISM as follows:

$$P_{min=?}[\mathbf{G}^{\leq N} (pollution \geq Y \Rightarrow P_{\geq 1}[\mathbf{F}^{\leq K} (pollution \leq X)])]$$

4 Methodological Guidelines

In this section, we provide methodological guidance to support the design and implementation of resilient STSs. We provide guidance from two perspectives:

- From the perspective of **STS designers**. Here, the challenge is to ensure that the STS meets the specified requirements while preserving agent autonomy.
- From the perspective of **stakeholders** specifying requirements for the overall system. Here, the challenge is to ensure that the requirements are *feasible*, i.e., that it is not possible to design an STS that satisfies the given requirements.

These are deliberately framed as *methodological guidelines* as opposed to a procedure because there is significant human input and judgment involved. There are multiple options available to the **STS designer** in terms of how the constituent agents (and in particular available agent actions) are designed and in terms of the kinds of norms that should be specified. There are, similarly, multiple options available to stakeholders in terms of how they might design requirements.

967 4.1 The STS Designer Perspective

968 The steps below serve as a guide for the STS designer to develop an initial
969 version of the STS specification and conduct an analysis to ensure stakeholder
970 requirements are satisfied. If the stated stakeholder requirements are not satis-
971 fied, the STS designer can iterate the specification to refine it and ensure that
972 stakeholder requirements are satisfied.

973 *Step1: Defining the agents constituting the STS.* The STS designer needs to
974 define each constituent agent as a collection of actions, based on the known
975 capabilities of each agent.

976 *Step2: Defining the augmented probabilistic state transition model.* In gen-
977 eral, a probabilistic state transition model describes the likelihood of
978 achieving a resultant state if a specific action is performed in a given
979 state. These probabilities are either provided as input (in the case of
980 *transition probability*) or computed from an agent's willingness to violate
981 a norm (the parameter w which is used in computing the *selection prob-
982 ability* for an action). In this step, the **STS designer** needs to create
983 an augmented probabilistic state transition model that serves as the
984 model for model-checking PRISM based on the two sets of probabilities:
985 *selection probabilities* and *transition probabilities*.

986 *Step3: Defining the norms governing the STS.* There are some common prin-
987 ciples guiding the specification of the initial set of norms given a set of
988 available agents and a set of requirements. A resilience requirement spec-
989 ifies two kinds of conditions: (1) Conditions that flag a departure from
990 the "normal" operating mode and (2) Conditions that flag the restora-
991 tion of the normal operating mode. The overall requirement is augmented
992 with a deadline (in terms of the number of steps) within which normal
993 operations are restored, and a lower bound on the probability that this
994 will happen. The normal operating mode can be achieved via a combina-
995 tion of commitments and prohibitions (e.g., *CompanyNear* committing to
996 *Hospital* to produce at least 100 PPE units per day, but being prohibited
997 by *Regulator* from polluting at a rate higher than 60 parts per million
998 per day). The deviation from the normal operating mode is flagged via
999 a condition where the minimum PPE production target is not met and
1000 the maximum permitted pollution level is exceeded. Achievement and
1001 maintenance requirements can similarly be specified via combinations of
1002 commitments and prohibitions.

1003 *Step4: Using PRISM to determine requirements satisfaction.* Our technique
1004 enables the encoding of requirements in the form of PCTL properties that
1005 can be provided to PRISM as input, and an augmented probabilistic state
1006 transition model. When we run PRISM with these inputs, we obtain as
1007 output an indication of whether the PCTL properties (and hence the orig-
1008 inal requirements) are satisfied. If PRISM indicates that these properties
1009 have not been satisfied, it triggers a redesign of the STS.

1010 *Step5: Redesigning the STS in light of model checking results.* Redesigning
1011 the STS can involve (1) Adding or removing agent actions and (2)

Adding or removing norms. Adding an action, if feasible, is one approach to resolving a case of requirements failure. If the analysis using PRISM reveals that an achievement goal/requirement is not being satisfied, a new action (for instance to produce PPE at a faster rate) can serve as a potential (re-design) solution. A violation of a maximum permissible pollution requirement can similarly be resolved by disallowing/removing an action that permits a faster rate of PPE production, which comes with a concomitant higher rate of pollution. Adding a new norm in the form of a commitment to produce higher levels of PPE can help resolve a situation where an achievement goal/requirement involving meeting a production target for PPE is not being satisfied.

4.2 The Stakeholder Perspective

The stakeholder may be required to engage in an iterative process of progressively relaxing the applicable requirements if the STS designer determines that revising the STS specification (Step 5 in Section 4.1) cannot meet the requirements.

Modification of the applicable requirements can involve three kinds of changes:

- Modifying the associated conditions: In our examples, these mainly involve inequalities on state variables (such as the number of PPE units, or the pollution level). These are entirely left to human judgment.
- Modifying the number of steps (i.e., the deadline) and the minimum associated probability: Both of these can be conjointly modified and represent a trade-off space for the stakeholder. Consider, for instance, Figure 3 which shows the minimum probabilities achievable for a given number of steps. Figure 6 similarly illustrates the trade-off space in the context of a resilience requirement. Accessing a visualization of the trade-off space will make it easier for a stakeholder to decide on an appropriate relaxation of a requirement.

5 Prototype Implementation and Requirement Verification Results

This section describes the evaluation we carried out, for our approach. We describe the requirement verification results of our use case. Then, we discuss our experimental settings, scalability of the implementation, performance measures, and finally, report our results.

To construct and analyse a model with PRISM, it must be specified in the PRISM language, a simple, state-based language. Specifically, the PRISM language is composed of modules and variables. The values of these variables at any given time constitute the state of the module. The behavior of the module is described by a set of commands.

$[action]guard \Rightarrow prob_1 : update_1 + \dots + prob_n : update_n$

1059 The **guard** is a predicate over all the variables in the model. Each **update**
1060 describes a transition that can be made by a module where the guard is true. A
1061 transition relationship of the system is specified by allocating new values to the
1062 variables in the module. Each update is also assigned a **probability** (or in some
1063 cases a rate) - assigned to the corresponding transition. If a module has more
1064 than one variable, then updates describe the new value for each one of them.
1065 The initial state of a model is then defined by the initial value of all variables
1066 (see Appendix A.1) for the initial state in the PRISM model). However, the
1067 PRISM input language does not support external functions, such as function
1068 $prob(a_k, s, A, N)$ in Equation 1, to dynamically construct a state-transition
1069 model. Therefore, we have connected to PRISM programmatically to create a
1070 model building on the machinery described in Equation 1 (see Appendix A.2).
1071 PRISM provides a Java-based interface or API that allows users to interact
1072 with PRISM programmatically. This enables us to automate the construction,
1073 solving, and verification of the PRISM model. Algorithm 1 has been used to
1074 automatically construct the PRISM model by translating an STS specification
1075 into a PRISM model that incorporates dynamic state transitions. That is,
1076 state transition probabilities are recalculated after each transition depending
1077 on the values of the state variables. We show the snippets for state transitions
1078 in the PRISM model in Appendix A.3.

1079 The first step in our framework in Figure 1 is to capture the STS specifici-
1080 cation. Here, Listing 1 describes the STS specification for PPE manufacturing
1081 with two companies. Each company has three actions where two actions are
1082 used to manufacture PPE and one action describes a *null* action. We intro-
1083 duce a null action where the company does not want to manufacture PPE.
1084 Specifically, actions a_{10} and a_{20} are considered as null actions for Compa-
1085 nyNear and CompanyFar, respectively. The second step in our framework,
1086 stakeholder provide a set of requirements. So we use the requirements intro-
1087 duced in Section 3.4. Then, the STS specification is automatically converted
1088 into a PRISM model using Algorithm 1. Finally, we demonstrate the verifica-
1089 tion results for the requirements introduced in Section 3.4 to see how likely
1090 STS satisfies or fails each requirement. Section 5.1 and Section 5.3 have been
1091 used to show that if a requirement is not satisfied then the STS designer can
1092 revise the STS specification. Section 5.2 and Section 5.4 have been used to
1093 show that if a requirement is not possible to satisfy then the stakeholders can
1094 relax the requirement to satisfy with a given STS.

1095 Table 2 and Table 3 illustrate the probability needed to satisfy a given
1096 requirement with a given STS specification. If a requirement is noncompliant
1097 with an STS specification, then the STS specification needs to be refined to
1098 satisfy a given requirement. Table 2 and Table 3 also capture the results where
1099 we design the new requirements by relaxing the value of state variables based
1100 on results from the analysis phase. The results shown in Table 2 and Table 3
1101 have been derived by using methodological guidelines mentioned in Section 4.

1102 We present various cases, each corresponding to an STS specification, and
1103 verify whether the specification satisfies the requirement. In the paper, we
1104

have provided listings to show an STS specification for each case. However, the data such as outcomes of actions and execution probabilities used to conduct experiments is provided in the supplementary material [17].

5.1 Handling Strict Achievement and Maintenance Requirements

The requirement below is an instance of the achievement and maintenance requirement discussed above. The minimum probability that the company manufactures more than 200 PPE units within six timesteps and maintains the pollution level below the threshold of 70 ppm within six timesteps is 0.75. If we achieve a minimum probability of 0.75, then this requirement is satisfied. Table 2 shows the experimental settings for achievement and maintenance requirements. Each case in Table 2 corresponds to an STS specification.

We verify whether a case satisfies the requirement.

$$P_{min=?}[(pollution \leq 70) \quad U^{\leq 6} \quad (PPE \geq 200)] \quad (4)$$

Requirement	Case being verified	STS	Outcome
Equation 4	Original (Case 1)	Listing 1	Not satisfied even if the requirement is relaxed
	CompanyNear is productive (Case 2)	Listing 2	Satisfied if relaxed to 8 steps
	Both are productive (Case 3)	Listing 3	Satisfied as stated
Equation 5	Original	Listing 1	Satisfied as stated

Table 2 Experimental settings for achievement and maintenance.

Case 1 The STS in this case follows the specification of Listing 1. This requirement is only satisfied if both companies produce more than 200 PPE within six timesteps with a minimum probability of 0.75. The minimum probability that both companies produce more than 200 PPE within six timesteps is 0.69, hence, this requirement is not satisfied.

Figure 3 illustrates that the minimum probability to be satisfied for this property varies with the number of timesteps (k). Moreover, 0.73 is the probability achieved after increasing the k number of time steps to ten. Hence, this requirement cannot be satisfied within the timesteps allowed.

Case 2 The STS is revised for CompanyNear. The Listing 2 shows a revised specification of Listing 1 – commitment C_1 is replaced by C_3 (CompanyNear, Hospital, true, $PPE \geq 200$). After replacing commitment C_1 by C_3 , CompanyNear is more productive in terms of manufacturing more PPEs. Commitment

1151 **Listing 2** Revised STS specification of Listing 1 for CompanyNear to ensure achievement
1152 and Maintenance.

```
1153 1 Addition of norms or actions are shown in blue and
1154     deleted lines are shown with red strikethrough.
1155 2
1156 3 P1(CompanyNear, Regulator, true, pollution $\geq$ 60)
1157 4 C1(CompanyNear, Hospital, true, PPE $\geq$ 100)
1158 5 C3(CompanyNear, Hospital, true, PPE $\geq$ 200)
1159 6 P2(CompanyFar, Regulator, true, pollution $\geq$ 80)
1160 7 C2(CompanyFar, Hospital, true, PPE $\geq$ 100)
1161 8
1162 9 a11: m(true, {PPE, pollution}, {PPE+=[50, 100],
1163     pollution+=PPE*[0.2, 0.4]})  

1164 10 a12: m(true, {PPE, pollution}, {PPE+=[80, 120],
1165     pollution+=PPE*[0.5, 0.7]})  

1166 11 a10: m(true, {PPE, pollution}, {PPE+=[0, 0],
1167     pollution+=PPE*[0.0, 0.0]})  

1168 12
1169 13 a21: m(true, {PPE, pollution}, {PPE+=[50, 100],
1170     pollution+=PPE*[0.2, 0.4]})  

1171 14 a22: m(true, {PPE, pollution}, {PPE+=[80, 120],
1172     pollution+=PPE*[0.5, 0.7]})  

1173 15 a20: m(true, {PPE, pollution}, {PPE+=[0, 0],
1174     pollution+=PPE*[0.0, 0.0]})
```

1175

1176 C_3 enables the selection of an action for increasing the production of PPEs. The
1177 minimum probability achieved is 0.69 (achieved in six timesteps). Hence, this
1178 requirement is not satisfied. Figure 3 shows that the requirement is satisfied
1179 within eight timesteps with 0.77 minimum probability.
1180

1181 **Case 3** As we have seen in Case 2 this requirement is satisfied in eight or more
1182 timesteps and cannot be satisfied within six timesteps, hence, we revise STS
1183 for CompanyFar. STS follows Listing 3, a revision of Listing 2—commitment
1184 C_2 is replaced by $C_4(\text{CompanyFar}, \text{Hospital}, \text{true}, \text{PPE} \geq 200)$. In this case, the
1185 commitment to manufacture PPE has been increased for both the companies.
1186 Therefore, both companies are productive as they choose the actions that
1187 manufacture more PPE. Figure 3 shows the minimum probability achieved is
1188 0.75 (achieved in six timesteps). Hence, this requirement is satisfied.
1189

1190 It is interesting to note that the minimum probability of satisfying this
1191 requirement is less in Case 2 than in Case 1 with less than six timesteps. Here,
1192 CompanyFar is committed to producing more than 200 PPE, which leads it to
1193 select action a21 with priority, which results in high pollution in the short term.
1194 So, CompanyFar meets the social objective but the probability of choosing
1195 action a21 is expensive in terms of increasing pollution.
1196

Listing 3 Revised STS specification of Listing 2 to ensure achievement and maintenance.

```
1 Addition of norms or actions are shown in blue and  
2 deleted lines are shown with red strikethrough.  
3 P1(CompanyNear, Regulator, true, pollution $\geq$ 60)  
4 C3(CompanyNear, Hospital, true, PPE $\geq$ 200)  
5 P2(CompanyFar, Regulator, true, pollution $\geq$ 80)  
6 C2(CompanyFar, Hospital, true, PPE $\geq$ 100)  
7 C4(CompanyFar, Hospital, true, PPE $\geq$ 200)  
8  
9 a11: m(true, {PPE, pollution}, {PPE+=[50, 100],  
10 pollution+=PPE*[0.2, 0.4]})  
11 a12: m(true, {PPE, pollution}, {PPE+=[80, 120],  
12 pollution+=PPE*[0.5, 0.7]})  
13 a10: m(true, {PPE, pollution}, {PPE+=[0, 0],  
14 pollution+=PPE*[0.0, 0.0]})  
15 a21: m(true, {PPE, pollution}, {PPE+=[50, 100],  
16 pollution+=PPE*[0.2, 0.4]})  
17 a22: m(true, {PPE, pollution}, {PPE+=[80, 120],  
18 pollution+=PPE*[0.5, 0.7]})  
19 a20: m(true, {PPE, pollution}, {PPE+=[0, 0],  
20 pollution+=PPE*[0.0, 0.0]})
```

5.2 Handling Relaxable Achievement and Maintenance Requirements

As Figure 1 (the RENO framework) shows, an STS designer can analyze the various state variables to satisfy the requirements. The STS in this case follows the specification of Listing 1. After revising the STS specification (as mentioned in Case 3), we have seen that the minimum probability of achieving more than 200 PPE and maintaining the pollution below 70 ppm is 0.81. The STS designer may want to determine which values of PPE and pollution can satisfy the above property with probability 1.

Figure 4 and Figure 5 show these minimum probabilities when PPE and pollution vary for a range of values and the number of steps K. Figure 4 describes the minimum probabilities achieved when pollution is less than 70 ppm and PPE varies from 140 to 200. The result shows that the minimum probability 1 is achieved when we have PPE greater than 160 within 10 timesteps or PPE greater than 140 within eight timesteps.

Figure 5 indicates the minimum probabilities achieved when PPE exceeds 200 ppm and pollution varies from 70 ppm to 90 ppm. The result shows that the minimum probability 1 is achieved only when pollution is less than 90 ppm within 12 timesteps.

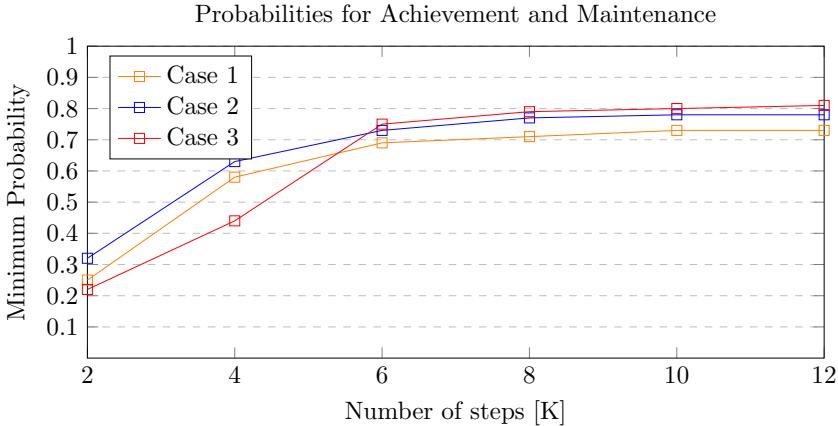


Fig. 3 Probabilities for achievement and maintenance.

These values of PPE and pollution indicate to the STS designer to provide new requirements by relaxing the value of PPE and pollution. The observations from Figure 4 and Figure 5 provide valuable insights for the STS designer. A minimum probability of 0.97 is achieved when we have PPE greater than 100 within six timesteps. Similarly, when pollution is less than 90 within six timesteps, a minimum probability of 0.92 is achieved. Based on these findings, the STS designer can express the new requirement as follows using these values of PPE and pollution.

$$P_{min=?}[(pollution \leq 90) \quad U^{\leq 6} \quad (PPE \geq 140)] \quad (5)$$

The minimum probability achieved is 1 (achieved in six timesteps). Hence, the requirement is satisfied with certainty (minimum probability of 1).

5.3 Handling Strict Resilience Requirements

Informally, we view resilience as an indicator of both the speed and likelihood with which a desirable property can be restored when the system of interest is found to be in a state that violates these properties.

This section describes the verification results of the resilience requirement. We aim to verify how the system reacts when pollution exceeds the given threshold value. The threshold value of the pollution represents a change in the operating environment of the system. Then, lowering the pollution level with probability 1 within a certain number of steps indicates the response to the system. Table 3 shows the experimental setting for resilience. Each case in Table 3 corresponds to an STS specification. We verify whether a case satisfies the requirement.

First, we would like to verify the minimum probability for a system entering a state where pollution is greater than 180 ppm and the system reduces the pollution level to below 100 ppm within five timesteps. This property can be written in PCTL as described in Equation 6.

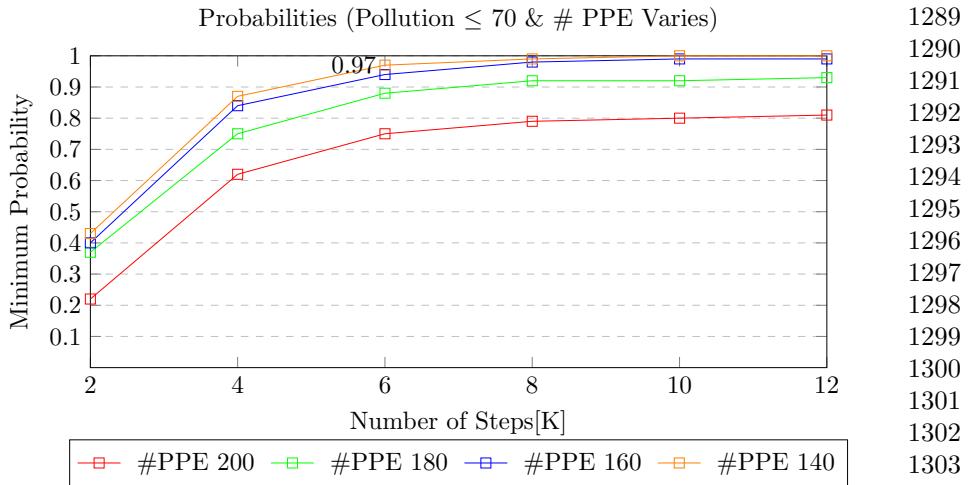


Fig. 4 Probabilities for achievement and maintenance when varying number of PPE units

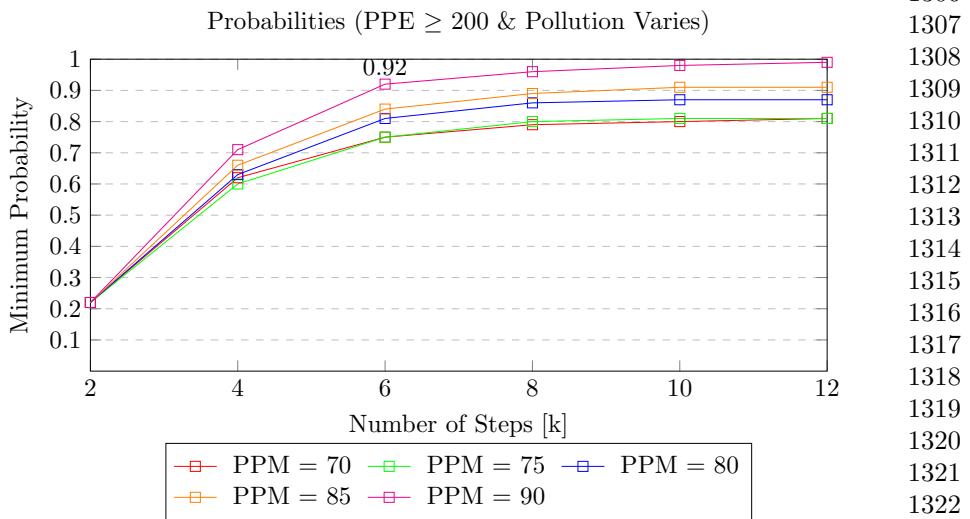


Fig. 5 Probabilities for achievement and maintenance when the varying amount of pollution (PPM).

$$filter(print, P_{min=?}[\mathbb{F}^{\leq 5} pollution \leq 100], pollution \geq 180) \quad (6)$$

Equation 6 is used to examine the minimum probability required for a state change from a state having pollution greater than 180 ppm to a state having pollution below 100 ppm within five timesteps. After verification, we get 1.0 as the minimum probability for this property.

1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334

Requirement	Case being verified	STS	Outcome
Equation 7	Original (Case 4)	Listing 1	Not satisfied even if the requirement is relaxed
	Both companies are productive (Case 5)	Listing 4	Satisfied as stated
Equation 8	Both are productive (Case 5)	Listing 4	Satisfied as stated

Table 3 Experimental settings for resilience.

Second, we are interested in ensuring that our system remains within the states that satisfy the above property (Equation 6) for seven timesteps with a minimum probability of 0.75. This property can be written in PCTL as follows:

$$P_{min=?}[\mathbb{G}^{\leq 7}(pollution \geq 180 \Rightarrow P \geq 1[\mathbb{F}^{\leq 5}(pollution \leq 100)])] \quad (7)$$

Equation 7 is a query for the minimum probability of the system staying in states where pollution goes above 180 ppm (threshold), the probability of lowering the pollution below 100 ppm in five timesteps with probability 1 for seven timesteps. Here, the property $P \geq 1[\mathbb{F}^{\leq 5}(pollution \leq 100)]$ represents the response that we are expecting from the system when there is a change in its operating environment ($pollution \geq 180$).

Case 4 The STS in this case follows the specification of Listing 1. The minimum probability achieved for the property of Equation 7 is 0.43. Hence, this requirement is not satisfied. Figure 6 shows that the minimum probability of recovering the system varies with the number of timesteps. When the number of steps increases, additional units of PPE are produced at the cost of increased pollution. It can be noticed from Figure 6 that the minimum probability achieved for the property of Equation 7 is 0.43 in seven timesteps, 0.3 in eight timesteps, and so on. Similarly, when the number of steps is decreased then fewer units of PPE are produced with the benefit of reduced pollution. There may be a possibility of requirement getting satisfied. Hence, we check the minimum probability achieved for the property of Equation 7 before seven timesteps. It can be noticed that even after reducing the number of steps, the property is not satisfied. For example, once we reduce the number of steps from seven to six and then to five the minimum probabilities achieved are 0.73 and 0.69, respectively. This property is not satisfied because none of the actions of the companies (Listing 1) reduce pollution adequately.

Case 5 We revise the STS specification (Listing 4) by adding commitment C_3 and C_4 for both companies. Now, both companies are committed to the Regulator to reduce the pollution level below 40 ppm if pollution increases above 90 ppm during the manufacturing of PPE. The revised STS specification

contains actions a13 and a23 for CompanyNear and CompanyFar respectively, to decrease pollution while manufacturing PPE. The STS in this case follows the specification (Listing 4) of Listing 1—commitments C_3 and C_4 are added for both companies.

Figure 6 shows that minimum probability achieved is 0.89 (achieved in seven timesteps). Hence, this requirement is satisfied. As we have seen in Case 4 with the original STS specification (Listing 1) the requirement is not satisfied even after we relax the number of steps but after revising STS specification (Listing 4) the property gets satisfied even though we increase the number of steps from seven to ten with 0.77 probability. After revising the STS specification (Listing 4), both companies meet the social objective of reducing pollution below 40 ppm whereas pollution above 90 ppm triggers actions that are less expensive in terms of environmental pollution. These newly introduced actions of each company manufacture the PPE and do not pollute the environment.

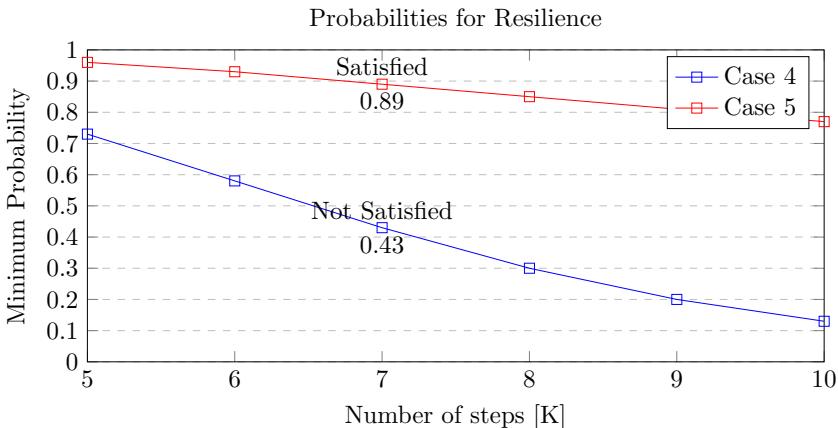


Fig. 6 Probabilities for resilience

5.4 Handling Relaxable Resilience Requirements

We now show how an STS designer can use our framework RENO in Figure 1 to design the new relaxed requirement. We use both the original STS specification (Listing 1) and the revised STS specification (Listing 4) to design the relaxed resilience requirement. The STS designer wants to explore which threshold value of pollution can satisfy the property (Equation 7) with a probability of 0.85.

Figures 7, 8, 9, and 10 show the probabilities calculated for the original STS specification (Listing 1) and the revised STS specification (Listing 4) when the pollution threshold varying for a range of values and number of steps.

Figures 7, 8, 9, and 10 indicate that the original STS specification (Listing 1) fails to achieve minimum probability 0.85 for the property of Equation 7

1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426

1427 Listing 4 Revised STS specification for CompanyNear and CompanyFar to ensure
 1428 resilience

```

1429 1 Addition of norms or actions are shown in blue.
1430 2
1431 3
1432 4 P1(CompanyNear, Regulator, true, pollution $\geq$ 60)
1433 5 C1(CompanyNear, Hospital, true, PPE $\geq$ 100)
1434 6 C3(CompanyNear, Regulator, pollution $\geq$ 90,
1435   pollution $\leq$ 40)
1436 7 P2(CompanyFar, Regulator, true, pollution $\geq$ 80)
1437 8 C2(CompanyFar, Hospital, true, PPE $\geq$ 100)
1438 9 C4(CompanyFar, Regulator, pollution $\geq$ 90,
1439   pollution $\leq$ 40)
1440 10
1441 11 a10: m(true, {PPE, pollution}, {PPE+=[0, 0],
1442   pollution+=PPE*[0.0, 0.0]})
1443 12 a11: m(true, {PPE, pollution}, {PPE+=[50, 100],
1444   pollution+=PPE*[0.2, 0.4]})
1445 13 a12: m(true, {PPE, pollution}, {PPE+=[80, 120],
1446   pollution+=PPE*[0.5, 0.7]})
1447 14 a13: m(true, {PPE, pollution}, {PPE+=[50, 100],
1448   pollution+=PPE*[-0.6, 0.0]})
1449 15
1450 16 a20: m(true, {PPE, pollution}, {PPE+=[0, 0],
1451   pollution+=PPE*[0.0, 0.0]})
1452 17 a21: m(true, {PPE, pollution}, {PPE+=[50, 100],
1453   pollution+=PPE*[0.2, 0.4]})
1454 18 a22: m(true, {PPE, pollution}, {PPE+=[80, 120],
1455   pollution+=PPE*[0.5, 0.7]})
1456 19 a23: m(true, {PPE, pollution}, {PPE+=[50, 100],
1457   pollution+=PPE*[-0.6, 0.0]})

1458
1459

```

1460 even after we increase the pollution threshold from 180 ppm to 240 ppm. For
 1461 instance, a minimum probability of 0.73 is achieved when the pollution thresh-
 1462 old is greater than 240 ppm within seven timesteps. However, Figure 7 shows
 1463 the minimum probability of 0.89 is achieved when the pollution threshold is
 1464 greater than 180 ppm within seven timesteps for the revised STS specification
 1465 (Listing 4). Similarly, minimum probabilities 0.93 (Figure 8), 0.95 (Figure 9),
 1466 and 0.98 (Figure 10) are achieved (within seven timesteps) when the pollution
 1467 threshold is greater than 200 ppm, 220 ppm, and 240 ppm, respectively.

1468 These threshold values of pollution provide an indication to the STS
 1469 designer to express a new requirement. Figure 10 illustrates that 90% of the
 1470 time, the system reduces the pollution below 100 ppm within five timesteps if
 1471 the pollution threshold value is greater than 240 ppm (change in the operat-
 1472 ing environment). It means that the system will remain in the state where the

pollution is below 100 ppm for the next seven, eight, nine, and ten timesteps with above 0.90 probability. Hence, the STS designer can design the new requirement for revised STS specification (Listing 4) as follows:

$$P_{min=?}[\mathbb{G} \leq 7(pollution \geq 240 \Rightarrow P \geq 1[\mathbb{F}^{\leq 5}(pollution \leq 100)])] \quad (8)$$

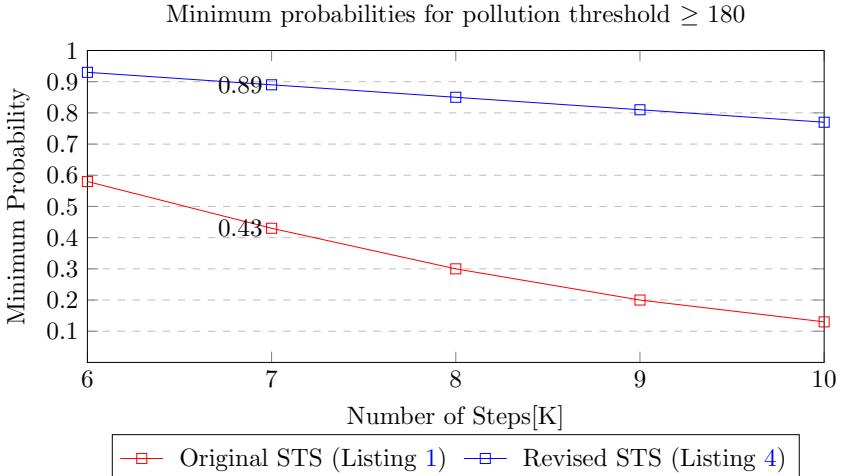


Fig. 7 Resilience (Equation 7) is not satisfied for the threshold value of pollution ≥ 180 when number of steps increased for original STS specification (Listing 1) but is satisfied for revised STS specification (Listing 4) within seven timesteps.

5.5 ReNo's Scalability

In this section, we will show computational cost of algorithms and scalability of ReNo. Section 5.5.1 has been used to show the time complexity of Algorithm 1. Section 5.5.2 has been used to show the time complexity of Algorithm 2. The scalability of ReNo has been explained in Section 5.5.3.

5.5.1 Analysis of time complexity for Algorithm 1

Theorem 2 *The time-complexity for the Algorithm 1 building PRISM model is $\mathcal{O}(|A|^2 * |S| * |N|)$.*

Proof From the initial state $s_0 \in S$, we consider each state in a set of states S . So, the outer loop executes at most $|S|$ times. Then, for each state, we consider each action in a set of actions A which means the inner loop executes at most $|A| * |S|$ times. Then, for each action, we compute the state transition probability by multiplying the action selection probability and action execution probability. The

1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518

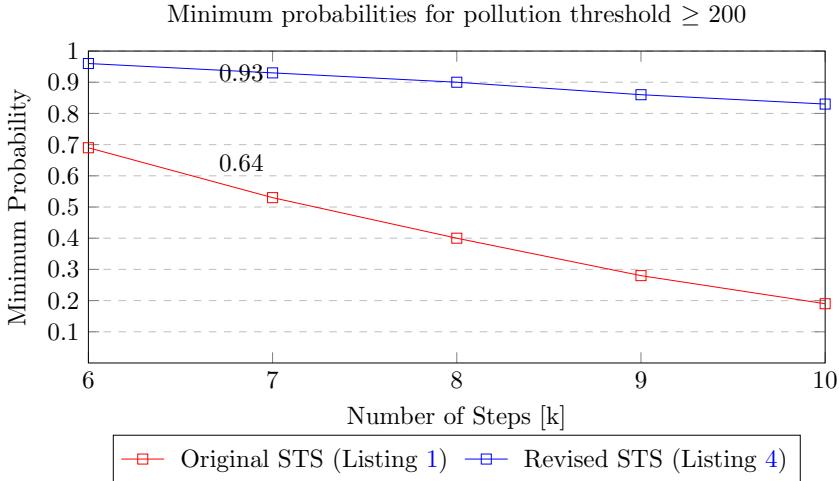


Fig. 8 Resilience (Equation 7) is not satisfied for the threshold value of pollution ≥ 200 when number of steps increased for original STS specification (Listing 1) but is satisfied for revised STS specification (Listing 4) within 7 timesteps.

time-complexity for Algorithm 2 computing selection probability in a given state s is $\mathcal{O}(|A| * |N|)$. Hence, Algorithm 2 executes at most $|A| * |S| * \mathcal{O}(|A| * |N|)$. Thus, the overall execution time is $\mathcal{O}(|A^2| * |S| * |N|)$.

In a worst-case scenario, the number of states in the augmented probabilistic state transition model representing the PRISM model can increase exponentially with the number of state variables and their possible assignments. For example, if there are v state variables with n possible assignments each, the total number of states becomes $2^{(v*n)}$, which grows exponentially as v or n increases.

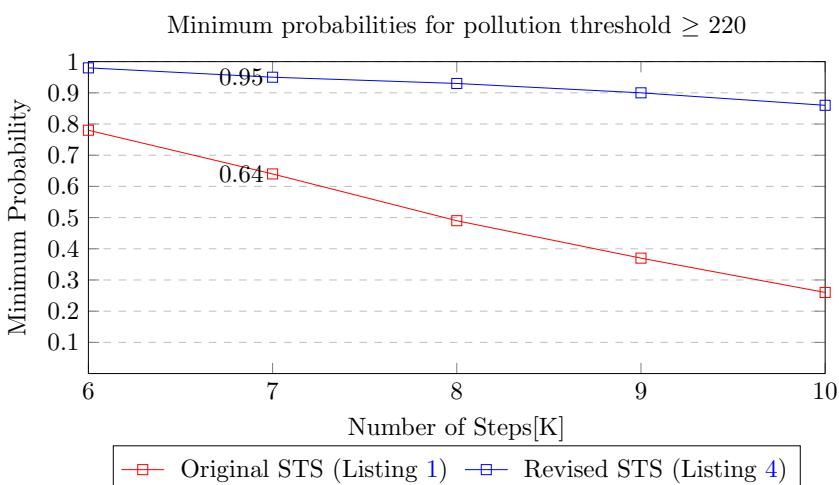


Fig. 9 Resilience (Equation 7) is not satisfied for the threshold pollution ≥ 220 when the number of steps increased for original STS specification (Listing 1) but is satisfied for revised STS specification (Listing 4) within 7 timesteps.

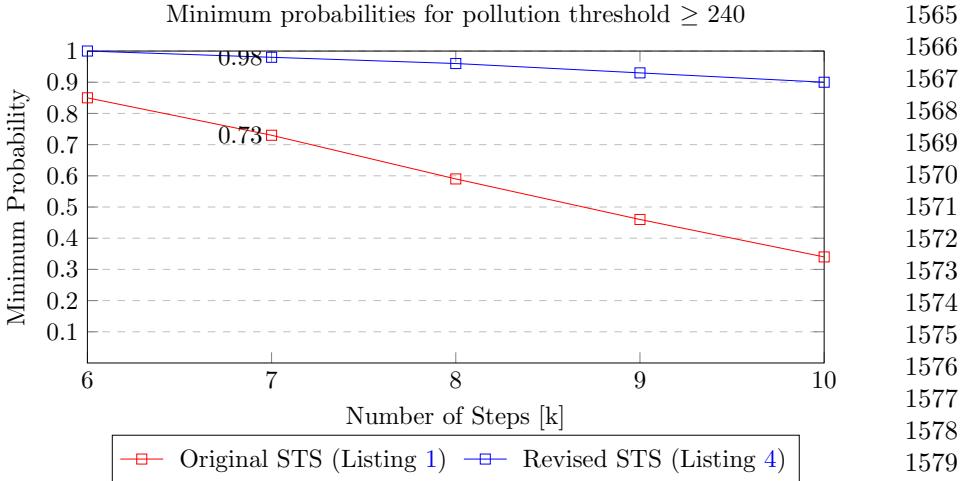


Fig. 10 Resilience (Equation 7) is not satisfied for the threshold value of pollution ≥ 180 when the number of steps increased for original STS specification (Listing 1) but is satisfied for revised STS specification (Listing 4) within seven timesteps.

However, in our work, we have utilized the PRISM model checker programmatically to construct a PRISM model from an STS specification. PRISM employs Multi-Terminal Binary Decision Diagrams (MTBDDs) to represent the state space efficiently. MTBDDs offer a compact and symbolic representation of sets of states, transitions, and properties, which facilitates efficient state space exploration and model checking. Consequently, although the time complexity of MTBDD-based techniques can be exponential in theory, their practical performance is often significantly improved through optimization strategies like symbolic representation and efficient state space exploration. As a result, the complexities derived from theoretical worst-case analysis can be overly pessimistic for MTBDD-based techniques. To gain insights into the actual model statistics and performance, we conducted experiments, which are detailed in Section 5.5.3.

□

5.5.2 Analysis of time complexity for Algorithm 2

Theorem 3 *The time-complexity for Algorithm 2 computing selection probability in a given state s is $\mathcal{O}(|A| * |N|)$.*

Proof In a given state s , we consider each action in a set of actions A . So, the outer loop executes at most $|A|$ times. Then for each action, we consider each norm in a set of norms N , which means the inner loop executes at most $|N| * |A|$ times. Thus, the overall execution time is $\mathcal{O}(|A| * |N|)$.

□

1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610

1611 **5.5.3 Model Statistics**

1612 In this section, we analyse the scalability of our RENO in terms of agents and
 1613 their actions. All experiments were executed on a machine running Windows
 1614 10 Enterprise edition with an Intel® Core™ i7-7700 CPU with a 3.60GHz
 1615 processor, 16.0 GB RAM, and a 64-bit Operating System. We performed all
 1616 experiments on the STS specification in Listing 1, where each agent has three
 1617 actions and three consequences except the null actions. We have used PRISM
 1618 version 4.6.1 to conduct experiments but any PRISM version can be used to
 1619 use our approach.

1620 Table 4 shows the statistics of PRISM models we have built for an increasing
 1621 number of agents and actions. Table 4 includes the number of agents and
 1622 the number of actions executed by all agents. For example, there are four
 1623 agents and twelve actions in Table 4 meaning that four agents all together
 1624 execute twelve actions.

1625 Table 4 also includes the number of states and transitions in the PRISM
 1626 model. The construction time is the time to build the augmented probabilistic
 1627 state transition model representing the PRISM model. The verification time
 1628 is the time taken to verify the system property written in PCTL.

1630
 1631 **Table 4** Model Statistics

Agents & Actions		Model		Time (s)	
#Agents	#Actions	States	Transitions	Construction time (s)	Verification time (s)
4	12	100394	961958	264.8	9.419
8	24	143266	1316691	451.649	6.504
12	36	147633	1331552	463.472	17.105
16	48	308512	2931077	1377.17	18.057
20	60	319938	3050028	1442.601	21.732
24	72	340511	3252277	1545.675	31.903

1642

1643

1644 Overall, the construction time is calculated by combining the following: (1)
 1645 the time taken by Algorithm 1 to generate a PRISM model from an STS specification,
 1646 (2) the time taken by Algorithm 2 to compute the selection probabilities
 1647 to select an action in the current state, (3) the time which is equivalent to the
 1648 time taken for the system description to be parsed and converted to Multi-
 1649 Terminal Binary Decision Diagrams (MTBDD) representing it, 4) the time to
 1650 perform reachability, identify the non-reachable states and filter the MTBDD
 1651 accordingly. We observe the state space explosion as the number of states of
 1652 a model grows exponentially in terms of the number of variables and parallel
 1653 components. Yet, they must be represented in a limited computer memory in
 1654 some form. For example, symbolic probabilistic model checking uses variants
 1655 of binary decision diagrams (BDD) to compactly represent the state spaces
 1656 of well-structured models in memory at the cost of verification runtime [25].

The complexity of PCTL model checking or requirement verification is linear in the size of a PCTL formula. (The size is defined as the number of logical connectives and temporal operators [25].)

However, as mentioned earlier, it is worth noting that the worst-case time complexity of building the PRISM model using Algorithm 1 can be exponential (see Section 5.5.1). Consequently, when dealing with large models, the computational time required for analysis may become excessively long. This can result in the impractical to complete the analysis within a reasonable timeframe.

6 Related Work

We review relevant approaches to RENO in the context of how to make a multiagent system more robust. Cheong and Winikoff [26] provide an approach to designing goal-oriented interactions that result in flexible and robust multiagent systems. In contrast, we define a sociotechnical system in computational terms as combining a social layer (characterized by norms) with a technical layer (characterized by actions and their effects). Our framework RENO provides an approach to make a socio-technical system (STS) more resilient. Socio-technical systems are an important concept in the study of systems of autonomous systems, as studied in the fields of multiagent systems and decentralized AI. The significance of sociotechnical systems arises from how they synthesize social and technical layers to develop a more complete understanding of how AI agents may be developed and deployed to solve problems of societal significance.

Chopra et al. [27] consider the robustness of an organization with respect to business contracts. For them, a contract is robust if it helps lead to behaviours that satisfy an organization’s goals and avoids undesirable outcomes. Their analysis is qualitative and concerns how an organization may be designed so it, for example, monitors and responds to potential or actual contract violations. In contrast, we are concerned with formal methods for evaluating a socio-technical system that includes multiple organizations.

More recent approaches [28, 29] have brought forth a notion of accountability in socio-technical systems. The idea is that, as stated above, the norms in an STS indicate what an agent may expect from another. Moreover, the agents can be held to accountable for any violations of the norms that may occur. An agent who violates a norm can be called upon to explain the violation or make corrections. Baldoni et al. [29] specify how responsibilities should be distributed in the STS. Chopra and Singh [28] show how capturing high-level requirements in terms of who is accountable for ensuring what can help create an STS that can better serve stakeholder needs by handling exception conditions more easily than systems where the requirements are at a low level. These approaches can benefit from the formal verification approach of this paper. If a multiagent system is verified with a certain probability to have the capability to recover before its execution, then this can be incorporated into the accountability mechanisms. Accountability mechanisms can be designed to

1703 monitor the performance of the system with respect to its requirements. For
1704 example, if the system fails to recover within the specified time frame, this can
1705 trigger an accountability mechanism that reports the failure to the relevant
1706 stakeholders and initiates a corrective action.

1707 In the literature, multiagent systems have addressed exception handling,
1708 but they view it primarily as a mechanism for dealing with specific errors
1709 that may arise in the system. Hæg [30] proposes a fault-handling approach
1710 for multiagent systems that introduce special agents called sentinels. Sentinels
1711 act as a control structure over the system and provide a layer of fault tol-
1712 erance. They do not participate in problem-solving but can intervene when
1713 necessary to choose alternative problem-solving methods. Similarly, Bungie
1714 [31] describes fault-tolerance patterns that can enhance the robustness of infor-
1715 mation protocols in a casual manner. Mandrake [32] is a programming model
1716 for decentralized applications that does not rely on infrastructure guarantees.
1717 It uses an information protocol that can be executed by agents in a shared-
1718 nothing environment using unreliable and unordered transport. Mandrake is
1719 focused on addressing faults that arise from violations of protocol expectations.
1720 In this context, fault handling can be seen as a form of exception handling,
1721 where agents with expectations can initiate recovery procedures to bring the
1722 system back into alignment with the protocol.

1723 Additionally, Baldoni et al. [29] propose the notion of accountability as a
1724 means of achieving robustness in multiagent systems. The concept of account-
1725 ability can aid in implementing effective exception-handling mechanisms.
1726 When an exceptional condition arises, the accountable agent(s) responsible for
1727 handling the situation is automatically notified with a corresponding account
1728 (i.e., exception) to take the necessary actions. Baldoni et al. [33] discuss
1729 how exception handling can be effectively integrated into distributed systems,
1730 which are implemented as multiagent systems, by utilizing the concept of
1731 responsibility at an organizational level. Gutierrez-Garcia's [34] approach in
1732 the context of normative multiagent systems involves modeling interaction
1733 protocols and exception handlers using obligations in deontic logic to handle
1734 exceptions.

1735 Since norms are crucial to the functioning of an STS, it is important to
1736 consider where they come from. Aydoğan et al. [35] address how stakeholders
1737 can negotiate about which norms to adopt in their STS based on the stake-
1738 holders' values. Values here refer to the key beliefs and desires of an entity,
1739 such as privacy, freedom, and safety. The negotiation is facilitated by ontology-
1740 based reasoning to promote their respective value. Incorporating stakeholder
1741 values means that the stakeholders will try harder to satisfy the norms of the
1742 STS than if their values were ignored. In their agent-based negotiation frame-
1743 work [35], an agent revises the STS specification by revising a set of norm
1744 revision rules. However, their approach does not provide formal reasoning for
1745 resilience of the sort developed in our paper. Our framework, RENO, offers
1746 a redesigned STS specification process that incorporates formal probabilis-
1747 tic reasoning. This enables us to satisfy stakeholders' requirements, including
1748

resilience, by adding or removing norms or by adding or removing agent actions. Our approach can potentially be enhanced to accommodate values by providing additional domain knowledge of how various actions performed by the agents promote or demote the values of the stakeholders of the STS and by including ontological reasoning about the application domain. 1749
1750
1751
1752
1753

D’Inverno [3] is relevant to this work, insofar as they define a special class of multiagent systems called *electronic institutions* which resemble institutions in human societies. Their approach doesn’t fully support autonomy in that a governor agent in an institution can overrule the decisions of a member agent. Additionally, their framework utilizes the state-based language Z to formally describe the behavior of a system in terms of its states and state transitions. State-based languages excel at capturing both the static structure and dynamic behavior of systems. However, to address the specific needs of probabilistic behaviors, our framework RENO incorporates a temporal logic language PCTL. PCTL enables the specification and reasoning of properties in systems with probabilistic behaviors. 1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764

Ajmeri et al.[36] focus on how norms emerge in a decentralized manner based on the interactions of the agents and how they give (positive or negative) sanctions to each other based on their apparent norm violations. In their approach, an agent who violates a norm can share its context with other agents as a weak explanation of why it violated the norm. This context can help the receiving agent decide if the norm violation was justified in the specific context, affecting its sanctioning of the violator. Ajmeri et al. show that the norms that emerge lead to improved outcomes for the participating agents. Agrawal et al. [37] address similar intuitions but for explicit norms. Although Ajmeri et al. and Agrawal et al. provide useful intuitions about the emergence of norms, they do not produce norms of complex logical forms. Although their approach enables an individual agent to evaluate a specific execution run based on its local observation, they do not address the evaluation of an STS as a whole. Potentially, their approach could be used as another tool for an STS designer by using simulations to create richer norms. Our approach could be used to evaluate the norms (i.e., the STS) for resilience. 1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780

Gasparini et al. [38] have used the contrary-to-duty structure to design the normative specification and have used a preference relation over possible worlds that captures levels of system robustness. A qualitative reward function has been used to check the levels of compliance. In RENO, MAS specification is defined as a social layer that includes norms and a technical layer that includes software components. Our framework uses probabilistic analysis to check resilience. Overall, Gasparini et al. address norm-driven multiagent planning in a probabilistic setting using Partially Observable Markov Decision Processes (POMDPs), and make provision for robustness analysis, but do not address the problems we pose. 1781
1782
1783
1784
1785
1786
1787
1788
1789
1790

Savarimuthu et al. [39–41] have provided an architecture to detect prohibition and obligation norms based on interactions between the agents. The 1791
1792
1793
1794

1795 prohibition and obligation detection algorithms utilizes association rule min-
1796 ing, a data mining technique, to identify sequences of events that may represent
1797 candidate norms. This architecture helps to modify or remove norms that
1798 are no longer relevant. Their work emphasises norm identification and has
1799 some relevance to our work. Our framework RENO assumes the existence
1800 of norms and shifts the attention to verifying the resilience of a STS. In
1801 our approach, we evaluate whether an STS can effectively recover and meet
1802 resilience requirements.

1803 Mahmoud et al. [42] propose a resource-aware adaptive punishment tech-
1804 nique that enables norm establishment with larger neighbourhood sizes than
1805 resource-unaware punishment. Their evaluation of the adaptive technique has
1806 been done via a simulation. Dell'Anna et al. [43] propose a mechanism to
1807 revise the norms automatically that consider agents' preferences. The norms
1808 are revised by revising associated sanctions at runtime. The relationship
1809 between the satisfaction or violation of the norms and the achievement of the
1810 system-level objectives is learned from system execution data using a Bayesian
1811 Network. Then considering the runtime knowledge and the agent's preference,
1812 they develop heuristic strategies that automatically revise the sanctions asso-
1813 ciated with the enforced norms. They evaluate their mechanism on a traffic
1814 simulator ring-road environment. In RENO, instead of focusing on sanctions,
1815 we examine commitment and prohibition norms. In terms of system properties,
1816 our focus is on probabilistic system properties, allowing for the consideration
1817 of uncertain system behavior. Within our framework, we have incorporated
1818 resilient requirements alongside achievement and maintenance properties to
1819 establish a resilient Socio-Technical System.

1820 Kafali et al. [11] propose a formal framework for the verification and refine-
1821 ment of the STS specification via social norms. They provide an agent-based
1822 simulation environment to mimic a social community consisting of multiple
1823 hospitals, physicians, and patients. They use an agent-based simulation to
1824 evaluate candidate STS designs and guide their refinement. A simulation envi-
1825 ronment is set up using parameters related to norms and mechanisms obtained
1826 from the STS specifications and requirements. Our approach is to design
1827 resilient sociotechnical systems by incorporating probabilistic model checking
1828 in a methodology for specifying a sociotechnical system and verifying time
1829 and quantity-constrained resilience requirements. We use probabilistic model
1830 checking to evaluate our approach.

1831 Cámera and De Lemos [44] provide an approach for verifying self-adaptive
1832 systems. The focus is on resilience properties, which assess the system's abil-
1833 ity to maintain reliable service provision despite environmental changes. They
1834 have considered Discrete-Time Markov Chains (DTMCs) as the probabilistic
1835 model employed to depict the system's behavior. The system's environment
1836 is stimulated for collecting execution traces to a build probabilistic model.
1837 PRISM is used to perform probabilistic model checking to evaluate the
1838 resilience of self-adaptive systems. Their framework has only relevance to
1839 our approach is verifying system property using probabilistic model checker
1840

PRISM. But our framework focuses on how to make resilient STS, by considering norms and actions. Model checking using PRISM combines probabilistic analysis and conventional reachability [45]. Unlike a simulation, our approach produces exact verification results of a property. Typically, there is a tradeoff: simulation gives us greater modeling realism, e.g., the ability to implement and experiment with various agent strategies, but it does not handle a formal notion of correctness. Model checking gives formal guarantees by checking properties against an STS model but requires a more abstract (simpler) model. Thus, model checking handles more limited phenomena, but with greater rigor.

Ostrom's notable contribution lies in her comprehensive focus on rules as the fundamental units of analysis in institutional theory and design. This emphasis highlights the vital role rules play in shaping institutional behavior and outcomes. Her work provides valuable insights into effective institutional design and management [46]. Social rules serve as the fundamental building blocks of institutional arrangements and are therefore essential elements in resilience analysis and design. They form the conceptual foundation upon which the analysis and design of resilient systems are built [46, 47]. Ostrom emphasizes that in the process of evolutionary institutional change, the variation in rules and norms is frequently a product of deliberate and rational design, rather than being driven solely by random factors. The work of Ostrom on the management of shared resources [48] is an example of how social regulation might provide a solution to problems such as the tragedy of the commons (the connection of such problems with normative multiagent systems has been made explicit in a number of papers, such as [49] and [50]). For instance, the authors [49] have used the perspective of self-governing institutions for common-pool resource (CPR) management, as defined by Ostrom, the concept of an institution encompasses the set of rules that outline the conditions governing the allocation and use of resources. These rules should be capable of being modified by additional rules, allowing for adaptation to the specific context in which the system operates. Additionally, the actions of individuals within the system can alter the environment, and external factors may also impact the system's dynamics [49]. Our framework RENO shares conceptual similarities to Ostrom's analysis and design of resilient systems, as outlined in her works ([46, 47]). The objective of our framework is to design and implement resilient socio-technical systems by adjusting either the social layer or the technical layer to fulfill the resilient requirements of stakeholders.

7 Conclusions and Future Directions

We propose RENO, a probabilistic framework for the design and verification of STSs involving social norms, technical actions, and probabilistic temporal stakeholder requirements. Our main contribution is to introduce the idea of socio-technical resilience and incorporate probabilistic model checking in an overall methodology for specifying STSs and verifying resilience requirements. Normative behaviour has been incorporated into MDPs by considering norms

1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886

1887 as explicit constraints on agents' transitions [51], or by expressing norms explic-
1888 itly as part of the state space and action space [52]. Other works on norm
1889 monitoring [53, 54] describe the expected normative behaviours of agents and
1890 resolved conflicts among norms using agents' objectives and preferences. We
1891 have incorporated norms explicitly into the PRISM model by calculating the
1892 probability of executing an action based on how the potential outcome of the
1893 action contributes towards the satisfaction or violation of the set of norms spec-
1894 ified in the STS. We have also provided an algorithm that translates a formal
1895 STS specification into a probabilistic model. The RENO framework provides
1896 a probabilistic guarantee that an STS meets its stakeholder requirements and
1897 recovers from failure within a certain period of time.

1898 RENO supports three core types of requirements: achievement and main-
1899 tenance requirements are essential to the norm literature (e.g., a commitment
1900 to achieve something to a certain level or a prohibition to maintain some-
1901 thing at a certain level). The third type is the resilience requirements (novel
1902 to RENO) to verify that an STS can recover from requirement failures. We
1903 believe these three types of requirements cover realistic verification scenarios,
1904 to help guide STS designers. Our findings suggest that RENO helps to build
1905 resilient STSs by demonstrating to the designer the tradeoffs between various
1906 STS specifications.

1907

1908 Future Work

1909

1910 If a requirement is not satisfied, then counterexamples can reveal details about
1911 why the requirement failed and what modifications can be made to the STS
1912 specification to satisfy the requirement. PRISM can produce a violating adver-
1913 sary for a property, for which the verification outcome is false. We will explore
1914 employing such adversaries to guide the designer towards a revised STS.

1915 Another important area of future development is the automation of the
1916 iterative process of STS re-design and requirements relaxation or modifica-
1917 tion based on the results of probabilistic model checking. Figuring out what
1918 relaxations are appropriate for stakeholders can involve sophisticated reason-
1919 ing about goal conflicts [23] as well as creativity [22, 55] more generally. We
1920 currently offer methodological guidelines for doing this (to be followed by the
1921 STS designer and by stakeholders who specify requirements), but there is con-
1922 siderable reliance on human judgment for key decisions such as which STS
1923 re-designs or which requirements relaxations will lead us to a faster resolution
1924 of the problem of an STS design not realizing the specified requirements.

1925

1926 Acknowledgments

1927

1928 We dedicate this article to the memory of Professor Aditya Ghose, who passed
1929 away unexpectedly in February 2023.

1930 This work was initiated with support from a grant from the University of
1931 Wollongong and NC State University through a collaboration network called
1932 the University Global Partnership Network.

MPS additionally thanks the US National Science Foundation (grant IIS-1908374) for partial support.	1933
	1934
	1935
	1936
	1937
	1938
	1939
	1940
	1941
	1942
	1943
	1944
	1945
	1946
	1947
	1948
	1949
	1950
	1951
	1952
	1953
	1954
	1955
	1956
	1957
	1958
	1959
	1960
	1961
	1962
	1963
	1964
	1965
	1966
	1967
	1968
	1969
	1970
	1971
	1972
	1973
	1974
	1975
	1976
	1977
	1978

1979 References

1980

1981 [1] Castelfranchi, C. Modelling social action for AI agents. *Artificial Intelligence* **103** (1–2), 157–182 (1998). [https://doi.org/10.1016/S0004-3702\(98\)00056-3](https://doi.org/10.1016/S0004-3702(98)00056-3) .

1984 [2] Norman, T. J. & Reed, C. A logic of delegation. *Artificial Intelligence* **174** (1), 51–71 (2010). <https://doi.org/10.1016/j.artint.2009.10.001> .

1987 [3] d'Inverno, M., Luck, M., Noriega, P., Rodríguez-Aguilar, J. A. & Sierra, C. Communicating open systems. *Artificial Intelligence* **186**, 38–94 (2012). <https://doi.org/10.1016/j.artint.2012.03.004> .

1991 [4] Singh, M. P. in *Group ability and structure* (eds Demazeau, Y. & Müller, J.-P.) *Decentralized Artificial Intelligence, Volume 2* 127–145 (Elsevier/North-Holland, Amsterdam, 1991). Revised proceedings of the *2nd European Workshop on Modeling Autonomous Agents in a Multi Agent World (MAAMAW)*, St. Quentin en Yvelines, France, August 1990.

1997 [5] Chopra, A. K. & Singh, M. P. *Sociotechnical systems and ethics in the large*, 48–53 (ACM, New Orleans, 2018).

2000 [6] Kampik, T. *et al.* Governance of autonomous agents on the web: Challenges and opportunities. *ACM Transactions on Internet Technology (TOIT)* **22** (4), 104:1–104:31 (2022). <https://doi.org/10.1145/3507910> .

2004 [7] Chopra, A. K. *et al.* *Protos: Foundations for engineering innovative sociotechnical systems*, 53–62 (IEEE Computer Society, Karlskrona, Sweden, 2014).

2008 [8] Jones, A. J. I., Artikis, A. & Pitt, J. The design of intelligent socio-technical systems. *Artificial Intelligence Review* **39** (1), 5–20 (2013). URL <http://dx.doi.org/10.1007/s10462-012-9387-2>. <https://doi.org/10.1007/s10462-012-9387-2> .

2013 [9] Dalpiaz, F., Giorgini, P. & Mylopoulos, J. Adaptive socio-technical systems: A requirements-based approach. *Requirements Engineering* **18** (1), 1–24 (2013). <https://doi.org/10.1007/S00766-011-0132-1> .

2016 [10] Kafali, Ö., Ajmeri, N. & Singh, M. P. Revani: Revising and verifying normative specifications for privacy. *IEEE Intelligent Systems (IS)* **31** (5), 8–15 (2016). <https://doi.org/10.1109/MIS.2016.89> .

2020 [11] Kafali, Ö., Ajmeri, N. & Singh, M. P. Specification of sociotechnical systems via patterns of regulation and control. *ACM Transactions on Software Engineering and Methodology (TOSEM)* **29** (1), 7:1–7:50 (2020). <https://doi.org/10.1145/3365664> .

[12] Singh, M. P. Norms as a basis for governing sociotechnical systems. *ACM Transactions on Intelligent Systems and Technology (TIST)* **5** (1), 21:1–21:23 (2013). <https://doi.org/10.1145/2542182.2542203> . 2025
2026
2027
2028

[13] Singh, A. M. & Singh, M. P. Wasabi: A conceptual model for trustworthy artificial intelligence. *IEEE Computer* **56** (2), 20–28 (2023). <https://doi.org/10.1109/MC.2022.3212022> . 2029
2030
2031
2032
2033
2034

[14] Kwiatkowska, M., Norman, G. & Parker, D. *PRISM: Probabilistic symbolic model checker*, 200–204 (Springer, London, 2002). 2035
2036
2037
2038

[15] Hansson, H. & Jonsson, B. A logic for reasoning about time and reliability. *Formal Aspects of Computing* **6** (5), 512–535 (1994). <https://doi.org/10.1007/BF01211866> . 2039
2040
2041
2042
2043

[16] Ciesinski, F. & Größer, M. in *On probabilistic computation tree logic* (eds Baier, C., Haverkort, B. R., Hermanns, H., Katoen, J.-P. & Siegle, M.) *Validation of Stochastic Systems* 147–188 (Springer, 2004). URL https://doi.org/10.1007/978-3-540-24611-4_5. 2044
2045
2046
2047

[17] Mahala, G., Kafali, O., Dam, H., Ghose, A. & Singh, M. P. Replication package (2023). URL <https://anonymous.4open.science/r/JAAMAS-ReNO-C6EE>. 2048
2049
2050
2051
2052

[18] Verhagen, H., Neumann, M. & Singh, M. P. Normative multi-agent systems: Foundations and history. *Handbook of normative multiagent systems. College Publications* 3–25 (2018). URL <http://www.collegepublications.co.uk/downloads/handbooks00004.pdf> . 2053
2054
2055
2056
2057

[19] Singh, A. M. & Singh, M. P. Norm deviation in multiagent systems: A foundation for responsible autonomy. *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)* 289–297 (2023). <https://doi.org/10.24963/ijcai.2023/33> . 2058
2059
2060
2061

[20] Liaskos, S., Khan, S. M. & Mylopoulos, J. Modeling and reasoning about uncertainty in goal models: A decision-theoretic approach. *Software and Systems Modeling* 1–24 (2022). <https://doi.org/10.1007/S10270-021-00968-W> . 2062
2063
2064

[21] Kwiatkowska, M., Norman, G. & Parker, D. *PRISM 4.0: Verification of probabilistic real-time systems*, 585–591 (Springer, Snowbird, Utah, 2011). 2065
2066
2067
2068
2069
2070

[22] Murukannaiah, P. K., Ajmeri, N. & Singh, M. P. Acquiring creative requirements from the crowd: Understanding the influences of personality and creative potential in crowd RE. *Proceedings of the 24th IEEE International Requirements Engineering Conference (RE)* 176–185 (2016). <https://doi.org/10.1109/RE.2016.68> .

2071 [23] Murukannaiah, P. K., Kalia, A. K., Telang, P. R. & Singh, M. P.
2072 Resolving goal conflicts via argumentation-based analysis of competing
2073 hypotheses. *Proceedings of the 23rd IEEE International Requirements
2074 Engineering Conference (RE)* 156–165 (2015). <https://doi.org/10.1109/RE.2015.7320418> .
2076

2077 [24] Singh, M. P. Norms as a basis for governing sociotechnical systems. *ACM
2078 Transactions on Intelligent Systems and Technology (TIST)* **5** (1), 1–23
2079 (2014). <https://doi.org/10.1145/2542182.2542203> .
2080

2081 [25] Parker, D. A. *Implementation of symbolic model checking for probabilistic
2082 systems*. Ph.D. thesis, University of Birmingham (2003).
2083

2084 [26] Cheong, C. & Winikoff, M. P. in *Hermes: Designing flexible and robust
2085 agent interactions* (ed. Dignum, V.) *Handbook of Research on Multi-Agent
2086 Systems: Semantics and Dynamics of Organizational Models* Ch. 5, 105–
2087 139 (IGI Global, Hershey, Pennsylvania, 2009). URL <https://doi.org/10.4018/978-1-60566-256-5.ch005>.
2088

2089 [27] Chopra, A. K. *et al.* Analyzing contract robustness through a model of
2090 commitments. *Proceedings of the 11th International Workshop on Agent
2091 Oriented Software Engineering (AOSE 2010)* (6788), 17–36 (2011). https://doi.org/10.1007/978-3-642-22636-6_2 .
2093

2094 [28] Chopra, A. K. & Singh, M. P. Accountability as a foundation for require-
2095 ments in sociotechnical systems. *IEEE Internet Computing* **25** (6), 33–41
2096 (2021). <https://doi.org/10.1109/MIC.2021.3106835> .
2097

2098 [29] Baldoni, M., Baroglio, C., Micalizio, R. & Tedeschi, S. Accountability
2099 in multi-agent organizations: from conceptual design to agent program-
2100 ming. *Autonomous Agents and Multi-Agent Systems* **37** (1), 7 (2023).
2101 URL <https://doi.org/10.1007/s10458-022-09590-6>. <https://doi.org/10.1007/s10458-022-09590-6> .
2102

2103

2104 [30] Hägg, S. A sentinel approach to fault handling in multi-agent systems.
2105 *Multi-Agent Systems Methodologies and Applications: Second Australian
2106 Workshop on Distributed Artificial Intelligence Cairns, QLD, Australia,
2107 August 27, 1996 Selected Papers 2* 181–195 (1997). <https://doi.org/10.1007/BF0030090> .
2108

2109

2110 [31] Christie V, S. H., Chopra, A. K. & Singh, M. P. Bungie: Improving
2111 fault tolerance via extensible application-level protocols. *IEEE Computer*
2112 **54** (5), 44–53 (2021). <https://doi.org/10.1109/MC.2021.3052147> .
2113

2114 [32] Christie, S. H., Chopra, A. K. & Singh, M. P. Mandrake: multiagent
2115 systems as a basis for programming fault-tolerant decentralized applica-
2116 tions. *Autonomous Agents and Multi-Agent Systems* **36** (1), 16 (2022).
2116

<https://doi.org/10.1007/s10458-021-09540-8> . 2117
2118

[33] Baldoni, M., Baroglio, C., Micalizio, R. & Tedeschi, S. Exception handling as a social concern. *IEEE Internet Computing* **26** (6), 33–40 (2022). URL <https://doi.org/10.1109/MIC.2022.3216272>. <https://doi.org/10.1109/MIC.2022.3216272> . 2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162

[34] Gutierrez-Garcia, J. O., Koning, J.-L. & Ramos-Corchado, F. F. An obligation approach for exception handling in interaction protocols. *2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology* **3**, 497–500 (2009). <https://doi.org/10.1109/WI-IAT.2009.334> .

[35] Aydoğan, R., Kafalı, Ö., Arslan, F., Jonker, C. M. & Singh, M. P. NOVA: Value-based negotiation of norms. *ACM Transactions on Intelligent Systems and Technology (TIST)* **12** (4), 45:1–45:29 (2021). <https://doi.org/10.1145/3465054> .

[36] Ajmeri, N., Guo, H., Murukannaiah, P. K. & Singh, M. P. Robust norm emergence by revealing and reasoning about context: Socially intelligent agents for enhancing privacy. *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)* 28–34 (2018). <https://doi.org/10.24963/ijcai.2018/4> .

[37] Agrawal, R., Ajmeri, N. & Singh, M. P. Socially intelligent genetic agents for the emergence of explicit norms. *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)* 10–16 (2022). <https://doi.org/10.24963/ijcai.2022/2> .

[38] Gasparini, L., Norman, T. J. & Kollingbaum, M. J. Severity-sensitive norm-governed multi-agent planning. *Autonomous Agents and Multi-Agent Systems* **32** (1), 26–58 (2018). <https://doi.org/10.1007/S10458-017-9372-X> .

[39] Savarimuthu, B. T. R., Cranefield, S., Purvis, M. A. & Purvis, M. K. Obligation norm identification in agent societies. *Journal of Artificial Societies and Social Simulation* **13** (4), 3 (2010). <https://doi.org/10.18564/JASSS.1659> .

[40] Savarimuthu, B. T. R., Cranefield, S., Purvis, M. A. & Purvis, M. K. Identifying prohibition norms in agent societies. *Artificial intelligence and law* **21** (1), 1–46 (2013). <https://doi.org/10.1007/S10506-012-9126-7> .

[41] Savarimuthu, B. T. R. & Cranefield, S. Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent and Grid Systems* **7** (1), 21–54 (2011). <https://doi.org/10.3233/MGS-2011-0167> .

2163 [42] Mahmoud, S., Miles, S. & Luck, M. *Cooperation emergence under*
2164 *resource-constrained peer punishment*, 900–908 (International Foundation
2165 for Autonomous Agents and Multiagent Systems (IFAAMAS), 2016).
2166 URL <http://dl.acm.org/citation.cfm?id=2937056>.

2167

2168 [43] Dell'Anna, D., Dastani, M. & Dalpiaz, F. Runtime revision of sanctions
2169 in normative multi-agent systems. *Journal of Autonomous Agents and*
2170 *Multi-Agent Systems (JAAMAS)* **34** (2), 43.1–43.54 (2020). <https://doi.org/10.1007/s10458-020-09465-8> .

2171

2172 [44] Cámará, J. & De Lemos, R. Evaluation of resilience in self-adaptive sys-
2173 tems using probabilistic model-checking. *7th International Symposium on*
2174 *Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*
2175 53–62 (2012). <https://doi.org/10.1109/SEAMS.2012.6224391> .

2176

2177 [45] Kwiatkowska, M., Norman, G. & Parker, D. Quantitative analysis with
2178 the probabilistic model checker PRISM. *Electronic Notes in Theoretical*
2179 *Computer Science* **153** (2), 5–31 (2006) .

2180

2181 [46] Ostrom, E. *Understanding Institutional Diversity* (Princeton University
2182 Press, 2009).

2183

2184 [47] Ostrom, E. in *Developing a method for analyzing institutional change*
2185 (eds Batie, S. & Mercuro, N.) *Alternative Institutional Structures* 66–94
2186 (Routledge, 2008).

2187

2188 [48] Ostrom, E. A general framework for analyzing sustainability of social-
2189 ecological systems. *Science* **325** (5939), 419–422 (2009). <https://doi.org/10.1126/science.1172133> .

2190

2191

2192 [49] Pitt, J., Schaumeier, J. & Artikis, A. Axiomatization of socio-economic
2193 principles for self-organizing institutions: concepts, experiments and chal-
2194 lenges. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*
2195 **7** (4), 1–39 (2012). <https://doi.org/10.1145/2382570.2382575> .

2196

2197 [50] Bevan, C. *et al.* Factors in the emergence and sustainability of self-
2198 regulation. *Social Coordination: Principles, Artefacts and Theories*
2199 (2013). AISB Convention .

2200

2201 [51] Oh, J., Meneguzzi, F., Sycara, K. & Norman, T. J. An agent architecture
2202 for prognostic reasoning assistance. *Proceedings of the 22nd International*
2203 *Joint Conference on Artificial Intelligence* 2513–2518 (2011). <https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-418> .

2204

2205 [52] Fagundes, M. S., Ossowski, S., Luck, M. & Miles, S. Using normative
2206 Markov decision processes for evaluating electronic contracts. *AI Com-
2207 munications* **25** (1), 1–17 (2012). <https://doi.org/10.3233/AIC-2012-0511>

2208

2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254

[53] Modgil, S. *et al.* Monitoring compliance with e-contracts and norms. *Artificial Intelligence and Law* **23** (2), 161–196 (2015). <https://doi.org/10.1007/s10506-015-9167-9> .

[54] Dastani, M., Torroni, P. & Yorke-Smith, N. Monitoring norms: A multi-disciplinary perspective. *Knowledge Engineering Review* **33**, e25 (2018). <https://doi.org/10.1017/S026988918000267> .

[55] Murukannaiah, P. K., Ajmeri, N. & Singh, M. P. Enhancing creativity as innovation via asynchronous crowdwork. *Proceedings of the 14th ACM Web Science Conference (WebSci)* 66–74 (2022). <https://doi.org/10.1145/3501247.3531555> .

2255 **Appendix A Java Code: Translation of an STS
2256 specification into a PRISM model**
2257

2258 **A.1 Code Snippets for the Initial State in the PRISM
2259 Model**
2260

2261 The Listing 5 shows java code to return the initial state of the PRISM model.
2262 The list variables have the agent variable and all state variables. The set-
2263 Value(0, 1) in Line 2 in the Listing 5 shows that value 1 has been assigned to
2264 the variable at index 0 in variables list (here the variable agent is at index 0 in
2265 the list named variable). Similarly, random values have been assigned to other
2266 state variables.

2267 **Listing 5** Code Snippets for initial state in PRISM model construction

```
2268 1 public State getInitialState() throws  
2269     PrismException {  
2270 2 return new State(variables.size()).setValue(0,  
2271     3).setValue(1,20).setValue(2, 0); }  
2272  
2273
```

2274 **A.2 Code Snippets to Calculate the Action Selection
2275 Probability**
2276

2277 **Listing 6** Code Snippets to calculate action selection probability

```
2278 1 for(int m=0;m<AgentsMechanism.get(a).size();m++) {  
2279 2     if ((normType.equalsIgnoreCase("P"))) {  
2280 3         if (n.getValue().containsKey(varM)) {  
2281 4             int nConsequentValue =  
2282                 Integer.parseInt(n.getValue().get(AtomName));  
2283 5             double currentTarget = nConsequentValue -  
2284                 valueOfVariable.get(varM) ;  
2285 6             if (maxRange <= currentTarget) {  
2286 7                 prob = 1.0;  
2287 8 } else if (minRange >= currentTarget) {  
2288 9                 prob = 0.0;} else {  
228910                 prob = Double.parseDouble(new  
2290                     DecimalFormat("##.###").format((double)  
2291                     (currentTarget - minRange) / (maxRange -  
2292                     minRange));}  
229311                 double result = Double.parseDouble(new  
2294                     DecimalFormat("##.###").format(Math.pow(w,  
2295                     prob)));  
229612 }  
229713 }  
229814 }  
2299  
2300
```

A.3 Code Snippets for State Transitions in the PRISM Model

The Listing 7 shows the snippet of Java code of state transitions in the model. The function *computeTransitionTarget(int i, int offset)* is used to compute non-deterministic outcomes for executing all mechanisms of a selected agent. The Line 2 in the Listing 7 shows the current state. From Line 3 to Line 5 turns each agent using round-robin agent execution. In the initial state, the value of the agent variable is 1 (as explained in the use case) meaning that this agent is Company. In the next state, the value of the agent is 2 (i.e., Hospital); then the value of the agent is 3 (i.e., Regulator); and then the value of the agent is again 1 (i.e., Company). The remaining state variables are updated with new values in the Line 8 in the Listing 7.

The function *computeTransitionTarget(int i, int offset)* is used to compute non-deterministic outcomes of executing all mechanisms for a selected agent and the function *getTransitionProbability(int i, int offset)* is used to compute the state transition probability for each state transition.

Listing 7 Code Snippets for state transitions in PRISM model

```
1 public State computeTransitionTarget(int i, int
2     offset) throws PrismException {
3     State target = new State(exploreState) ;
4     if((agent==numberofAgents)&&(sum!=0) )  {
5         target.setValue(0,1);}
6     else if((sum!=0)) {
7         target.setValue(0,agent+1);}
8     int newTotalValue =
9         valueOfVariable.get(variables.get(p+1))
10    +Integer.parseInt(varValue[p]);
11    target.setValue(p+1, newTotalValue ) ;
12    return target;
13 }
14 public double getTransitionProbability(int i, int
15     offset) throws PrismException {
16     int agent = valueOfVariable.get(variables.get(0));
17     for(int t=0;
18         t<transitionProb.get(action).size();t++){
19         double probTra = transitionProb.get(action).get(t);
20         double prob = probTra * mechSelectionProb;
21         return prob ;}
```

2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346