Three-Valued Spatio-Temporal Logic: a further analysis on spatio-temporal properties of stochastic systems

Ludovica Luisa Vissat[1], Michele Loreti[2], Laura Nenzi[3], Jane Hillston[1], and Glenn Marion[4]

[1] School of Informatics, University of Edinburgh, UK
[2] DiSIA, University of Firenze, Italy
[3] DMG, University of Trieste, Italy
[4] Biomathematics and Statistics Scotland, Edinburgh, UK

Abstract. In this paper we present Three-Valued Spatio-Temporal Logic (TSTL), which enriches the available spatio-temporal analysis of properties expressed in Signal Spatio-Temporal Logic (SSTL), to give further insight into the dynamic behaviour of systems. Our novel analysis starts from the estimation of satisfaction probabilities of given SSTL properties and allows the analysis of their temporal and spatial evolution. Moreover, in our verification procedure, we use a three-valued approach to include the intrinsic and unavoidable uncertainty related to the simulation-based statistical evaluation of the estimates; this can be also used to assess the appropriate number of simulations to use depending on the analysis needs. We present the syntax and three-valued semantics of TSTL and a specific extended monitoring algorithm to check the validity of TSTL formulas. We conclude with two case studies that demonstrate how TSTL broadens the application of spatio-temporal logics in realistic scenarios, enabling analysis of threat monitoring and control programmes based on spatial stochastic population models.

1 Introduction

In many case studies, considering spatial structure is of key importance to better understand and predict the evolution of the system under study. For example, dispersive processes such as spread of disease, invasive species or fire spread have an intrinsic and fundamental spatial dimension that has to be included in the model. Spatial stochastic models provide a good representation of such system dynamics, typically studied through simulations. Correspondingly, the formal analysis of these spatial stochastic models has to also accommodate spatial and temporal modalities to be able to describe and verify properties about the spatio-temporal evolution of the specific systems.

Suitable analysis is provided by spatio-temporal logics and model checking. In most cases a statistical approach [1] is needed to estimate satisfaction probabilities of given properties, expressed using logical formulas. Simulation trajectories alone make it difficult to fully analyse the dynamic behaviour and to compare different systems, and the exhaustive exploration of all possible spatio-temporal trajectories is computationally infeasible. Using current simulation-based approaches the outcome is summary information about the satisfaction of properties over the spatial domain.

In our work we seek to add value to such information by providing a novel logic, called Three-Valued Spatio-Temporal Logic, to reason about spatial and temporal evolution of the satisfaction of these properties, giving further insight into the dynamic behaviour of the system under study. For example, in the analysis of the efficacy of a control measure for disease spread, we can verify whether the spread in a specific area will happen with probability under a given threshold over time. We can also identify the locations at highest risk, being surrounded by locations with high probability of becoming infected. The new TSTL atomic propositions are inequalities on the estimated satisfaction probabilities of given logical formulas (in the spatio-temporal logic SSTL [2] in this case, which formally describes and verifies properties of spatio-temporal trajectories), which are estimated using statistical model checking. This simulation-based evaluation has an intrinsic and unavoidable uncertainty, but frequently it is the only computationally feasible approach, requiring just an executable model. We use a three-valued approach to keep track of the associated uncertainty in the results of our model checking and we interpret the inequalities with different degrees of truth, using *true*, *false* and a third value *unknown*. This extension can be also used to give an indication of when more simulations are needed to make the evaluation of atomic propositions more precise and thus allowing stronger conclusions to be drawn. Conversely this enables initial explorations with relatively few simulations and assessment of whether they result in sufficient precision. We implemented the monitoring algorithms for the TSTL logical operators, to evaluate the satisfaction function of TSTL properties. The operators and the procedures are defined in a similar way to SSTL but on a different domain, dealing with three truth values.

Related work Several existing logics can be used to describe spatial properties of systems and estimate satisfaction probability values. Much of this work is based on topological models [3], looking at properties of subsets of points of topological spaces, whilst we take a more concrete representation of space. Other literature concerns spatial logics for process algebra with locations [4], used to study the mobility of concurrent systems; here space is represented as a tree and locations are nested. Based on a graph structure, there are logics such as the Multiprocess Network Logic [5], which can express spatio-temporal properties in discrete time. Considering stochastic systems, there are existing logics for expressing properties on probabilities, such as Probabilistic Computation Tree Logic (PCTL) [6] and Continuous Stochastic Logic (CSL) [7]. In these cases, the analysis is limited to temporal aspects, without spatial modalities, while our novel approach considers both. Three-valued logics, such as ours, with just one additional truth value, are a simple case in the field of multi-valued logics [8]. The initial concept was created by Łukasiewicz [9] and developed further by different logicians, such as Kleene [10], introducing the concept of "undefined" dealing with partial recursive functions. The three-valued approach is used in [11], for the definition of a new abstraction method for fully probabilistic systems and in [12], for model checking of Discrete-Time Markov Chains. We are not aware of any current use of a three-valued logic approach in the field of spatio-temporal analysis of stochastic systems.

Paper structure. The paper is structured as follows: Section 2 introduces notation and background work on SSTL while Section 3 presents the novel logic TSTL. Section 4 introduces the process algebra MELA we used to perform stochastic simulations, the monitor jSSTL we used to verify SSTL properties and how we linked all these aspects together to verify TSTL properties. Section 5 and 6 present two different case studies and applications of TSTL. Section 7 reports discussion and future directions for investigation while conclusions are reported in Section 8.

2 Background

In this section we introduce some fundamental concepts and notation that we will use in this paper aligned with the syntax and semantics of the existing spatio-temporal logic SSTL.

Notation We define a *spatial population model*, on a discrete representation of space; it describes a large number of different agents that can perform actions, take different states, interact and move between different locations. More formally, a *spatial population model* $\mathcal{M}$ is defined as a tuple $\mathcal{M} = (\mathcal{S}, G, \mathbf{X}, \mathbf{X}_0, Tr)$ where:

- $\mathcal{S} = \{1, \dots, n\}$ is the set of states that the population agents can take.
- $G = (L, E, w)$, a *finite weighted undirected graph* that represents the current choice of underlying spatial structure of the spatial population model:
  - L is the finite set of locations (nodes)
  - $E \subseteq L \times L$ is the set of connections (edges)
  - $w : E \to \mathbb{R}_{\geq 0}$ is the function *cost* (weights). We extend w to E^*, the transitive closure of E (set containing all the pairs of connected nodes). w gives the sum of costs of the shortest path between two different nodes, where this shortest path is the one that minimizes the sum of the costs.
- $\mathbf{X} : L \to \mathbb{R}^n$, where $\mathbf{X}(l) = (X_1, \dots, X_n) \in \mathbb{R}^n$ is the state vector, that represents the state of the population in each location. The entries of the vector $\mathbf{X}(l)$ represent the number of agents in location l in the i^{th} state; therefore, to be more specific, these *counting variables* are $X_i \in \mathbb{N}_0$.
- $\mathbf{X}_0 : L \to \mathbb{R}^n$, where $\mathbf{X}_0(l)$ is the initial state of the state vector, for each location.
- Tr is the set of transitions, $\tau_i = (\alpha_i, v_i, r_i)$, describing the events that change the global state of the system. Each transition consists of a label α_i in the label set $\mathcal{L}$, an update vector, $v_i : L \to \mathbb{R}^n$ recording the change to each counting variable in each location due to the transition, and a rate function r_i, which may depend on the global state of the system.

We can interpret the dynamical evolution of these models either stochastically as a Markov chain or deterministically as a system of Ordinary Differential Equations (ODEs); in this work we focus on stochastic spatio-temporal systems. We can describe the temporal evolution of our spatial population models using:

– σ, a *spatio-temporal trajectory* of $\mathcal{M}$. $\sigma : L \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ gives the state of the population vector for each location $l \in L$ and each time $t \in \mathbb{R}_{\geq 0}$
– Σ, a *set of spatio-temporal trajectories*, that will be used in the analysis.

SSTL Syntax Signal Spatio-Temporal Logic (SSTL) [2] is a spatial extension of Signal Temporal Logic (STL) [13], a temporal logic suitable for describing properties of real-valued signals. The syntax of SSTL is given by:

$$\varphi ::= \mu \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \, \mathcal{U}^{[t_1,t_2]} \, \varphi_2 \mid \; \diamondsuit_{[w_1,w_2]} \, \varphi \mid \varphi_1 \mathcal{S}_{[w_1,w_2]} \varphi_2$$

The SSTL *atomic proposition* μ is of the form $\mu \equiv (f \geq 0)$, $f : \mathbb{R}^n \rightarrow \mathbb{R}$, an inequality on expressions with population counts, given in the spatio-temporal trajectory. *Negation* $\neg$ and *disjunction* $\vee$ are the standard boolean operators and $\mathcal{U}$ is the *bounded until* operator. This temporal operator $\mathcal{U}$ is used to verify that the property φ_2 will be satisfied at some time instant in the interval $[t_1, t_2]$ and that at all preceding time instants φ_1 holds. SSTL introduces two spatial operators: the *bounded somewhere* operator $\diamondsuit_{[w_1,w_2]}$ and the *bounded surround* operator $\mathcal{S}_{[w_1,w_2]}$, with w_1, w_2 real values, $w_1 \leq w_2$. The *bounded somewhere* operator requires that the property φ holds in a location reachable from the current one, with a cost w, $w \in [w_1, w_2]$. The operator *bounded surround* describes the property of being surrounded by a φ_2-region, while being in a φ_1-region: the formula $\varphi_1 \mathcal{S}_{[w_1,w_2]} \varphi_2$ is true in a location l, if l belongs to a set of locations A where φ_1 holds, such that its external boundary $B^+(A)$ contains only locations satisfying φ_2. The external boundary of a subset of locations A is defined as $B^+(A) := \{l \in L \mid l \notin A \;\wedge\; \exists l' \in A \text{ s.t. } (l, l') \in E\}$. Moreover, the locations in the $B^+(A)$ have to be reached from location l with a cost w, $w \in [w_1, w_2]$. Examples of SSTL formulas will be provided throughout the paper.

SSTL Boolean semantics SSTL presents a *boolean* semantics that returns the value true/false ($\mathbb{B} = \{T, F\}$) depending on whether the observed trajectory satisfies the defined SSTL formula or not. The *boolean semantics* of a SSTL formula φ is interpreted over a spatio-temporal trajectory σ of $\mathcal{M}$, for each location $l \in L$ and at time $t \in \mathbb{R}_{\geq 0}$, given values in the set $\mathbb{B}$:

$$\beta(\mathcal{M}, \sigma, l, t, \varphi) \in \mathbb{B}$$

The satisfaction function β is defined as follows:

$$\begin{aligned}
\beta(\mathcal{M}, \sigma, l, t, \mu) &= \; \mu(\sigma(l, t)) \\
\beta(\mathcal{M}, \sigma, l, t, \neg\varphi) &= \; \neg\beta(\mathcal{M}, \sigma, l, t, \varphi) \\
\beta(\mathcal{M}, \sigma, l, t, \varphi_1 \vee \varphi_2) &= \; \beta(\mathcal{M}, \sigma, l, t, \varphi_1) \;\vee\; \beta(\mathcal{M}, \sigma, l, t, \varphi_2)
\end{aligned}$$

$$\beta(\mathcal{M},\sigma,l,t,\varphi_1\,\mathcal{U}^{[t_1,t_2]}\,\varphi_2) = \bigvee_{t'\in[t+t_1,t+t_2]} (\beta(\mathcal{M},\sigma,l,t',\varphi_2)\ \wedge \bigwedge_{t''\in[t,t')} \beta(\mathcal{M},\sigma,l,t'',\varphi_1))$$

$$\beta(\mathcal{M},\sigma,l,t,\diamondsuit_{[w_1,w_2]}\varphi) = \bigvee_{l'\in L,w(l,l')\in[w_1,w_2]} \beta(\mathcal{M},\sigma,l',t,\varphi)$$

$$\beta(\mathcal{M},\sigma,l,t,\varphi_1 S_{[w_1,w_2]}\varphi_2) = \bigvee_{A\in SR_l^{[w_1,w_2]}} (\bigwedge_{l'\in A} \beta(\mathcal{M},\sigma,l',t,\varphi_1)\ \wedge \bigwedge_{l''\in B^+(A)} \beta(\mathcal{M},\sigma,l'',t,\varphi_2))$$

where the surrounding region $SR_l^{[w_1,w_2]} = \{A \subseteq L \mid \forall l' \in A : 0 \leq w(l,l') \leq w_2\ \wedge\ \forall l'' \in B^+(A) : w_1 \leq w(l,l'') \leq w_2\}$.

Monitoring algorithms have been defined to evaluate the validity of SSTL properties, given a spatio-temporal trajectory, working inductively bottom-up on the parse tree of the formula. To make the verification procedure tractably computable, the time-domain has to be discretised, giving as output a piecewise constant approximation of the result. For this reason, in the analysis we talk about time-steps, although we start from discrete-event continuous-time simulations.

As discussed previously, in the study of stochastic systems we are generally interested in evaluating the probability that given properties are satisfied; a commonly used approach consists of estimating these values using statistical methods on a set of trajectories. Therefore, given a SSTL property φ, we shift the analysis from a single trajectory σ to a set of trajectories Σ, assigning to each trajectory a truth value, according to the boolean semantics. After this step we can estimate the satisfaction probability p^* of the formula φ, provided with a confidence interval. We define $\mathcal{P}_\beta$ over the set of trajectories Σ, in terms of β:

$$\mathcal{P}_\beta(\mathcal{M},\Sigma,l,t,\varphi) = (p^*,\delta) \tag{1}$$

where $(p^*,\delta) \in [0,1]\times[0,1]$ and represents the interval $[p^*-\delta, p^*+\delta]$.

$$p^* = \frac{|\Sigma_\top|}{|\Sigma|} \qquad \text{and} \qquad \delta = f_\delta(|\Sigma|,|\Sigma_\top|,\epsilon) \tag{2}$$

where $|\Sigma_T| = \{\sigma \in \Sigma \mid \beta(\mathcal{M},\sigma,l,t,\varphi) = T\}$ and δ is calculated with a given confidence level ϵ, according to a suitable function f_δ. There are several approaches to compute the confidence interval. For the sake of simplicity in our presentation we assume this interval to be symmetric. Given the boolean nature of the observations, SSTL uses the binomial proportion confidence interval and the most common choice for the calculation presupposes that the error distribution is approximated by a normal distribution. From this point on, all the results of SSTL monitoring are given at 95% confidence.

3 Three-Valued Spatio-Temporal Logic

We now present the novelty of our research, introducing the syntax and three-valued semantics of TSTL, providing also derived operators and a specific monitoring algorithm.

TSTL Syntax With the existing SSTL we are able to verify spatio-temporal properties of stochastic systems and estimate the satisfaction probabilities of given formulas. After this initial analysis we use our proposed extension to perform spatio-temporal analysis of these estimated values. The syntax of Three-Valued Spatio-Temporal Logic (TSTL) is given by:

$$\psi ::= \mathcal{P}_{<p}(\varphi) \mid \tilde{\neg}\psi \mid \psi_1 \tilde{\vee} \psi_2 \mid \psi_1 \tilde{\mathcal{U}}^{[t_1,t_2]} \psi_2 \mid \underset{[w_1,w_2]}{\tilde{\diamondsuit}} \psi \mid \psi_1 \underset{[w_1,w_2]}{\tilde{\mathcal{S}}} \psi_2$$

where $p \in [0,1]$ and φ is a given SSTL formula. The atomic TSTL formula $\mathcal{P}_{<p}(\varphi)$ expresses an inequality on the estimated satisfaction probability of the SSTL formula φ, checking if it is below the given threshold p. The logical TSTL operators link the TSTL propositions in a similar way to the SSTL ones, but working with estimated values and on a three-valued domain, as explained in the next section. We have *negation* $\tilde{\neg}$ and *disjunction* $\tilde{\vee}$ operators, *bounded until* $\tilde{\mathcal{U}}$, *bounded somewhere* $\underset{[w_1,w_2]}{\tilde{\diamondsuit}}$ and *bounded surround* $\underset{[w_1,w_2]}{\tilde{\mathcal{S}}}$. Conceptually all these operators are identical to the SSTL operators, but they operate on a different domain, reasoning about estimated satisfaction probabilities and not population counts. In the remainder we will show examples and differences between the two spatio-temporal logics; we will use the letter φ for SSTL formulas and ψ for TSTL ones.

Three-valued semantics TSTL presents a three-valued semantics that returns a truth values in $\mathbb{T} = \{T, U, F\}$ (true/unknown/false). The truth tables for TSTL negation $\tilde{\neg}$, disjunction $\tilde{\vee}$ and conjunction $\tilde{\wedge}$ (that can be defined in terms of negation and disjunction) are given by:

$\tilde{\neg}$	T	U	F
	F	U	T

$\tilde{\vee}$		ψ_2		
		T	U	F
ψ_1	T	T	T	T
	U	T	U	U
	F	T	U	F

$\tilde{\wedge}$		ψ_2		
		T	U	F
ψ_1	T	T	U	F
	U	U	U	F
	F	F	F	F

as for Kleene's logic of indeterminacy K3 [10]. The three-valued satisfaction function τ for the atomic TSTL proposition $\mathcal{P}_{<p}(\varphi)$ will return a value in $\mathbb{T}$:

$$\tau(\mathcal{M}, \Sigma, l, t, \mathcal{P}_{<p}(\varphi)) = [\![p^* <_\delta p]\!] \in \mathbb{T}$$

that is evaluated starting from the resulting (p^*, δ) given by $\mathcal{P}_\beta(\mathcal{M}, \Sigma, l, t, \varphi)$, as shown in the equations (1), (2). The associated truth value will be:

$$[\![p^* <_\delta p]\!] = \begin{cases} T & \text{if } p > p^* + \delta \\ U & \text{if } p \in [p^* - \delta, p^* + \delta] \\ F & \text{otherwise} \end{cases}$$

The three-valued satisfaction function τ for the TSTL operators is defined as follows, in an analogous manner as SSTL:

$$\tau(\mathcal{M}, \Sigma, l, t, \tilde{\neg}\psi) = \tilde{\neg}\tau(\mathcal{M}, \Sigma, l, t, \psi)$$

$$\tau(\mathcal{M}, \Sigma, l, t, \psi_1 \tilde{\vee} \psi_2) = \tau(\mathcal{M}, \Sigma, l, t, \psi_1) \tilde{\vee} \tau(\mathcal{M}, \Sigma, l, t, \psi_2)$$

$$\tau(\mathcal{M}, \Sigma, l, t, \psi_1 \tilde{\mathcal{U}}^{[t_1,t_2]} \psi_2) = \tilde{\bigvee}_{t' \in [t+t_1, t+t_2]} (\tau(\mathcal{M}, \Sigma, l, t', \psi_2) \tilde{\wedge} \tilde{\bigwedge}_{t'' \in [t,t')} \tau(\mathcal{M}, \Sigma, l, t'', \psi_1))$$

$$\tau(\mathcal{M}, \Sigma, l, t, \tilde{\diamondsuit}_{[w_1,w_2]} \psi) = \tilde{\bigvee}_{l' \in L, w(l,l') \in [w_1,w_2]} \tau(\mathcal{M}, \Sigma, l', t, \psi)$$

$$\tau(\mathcal{M}, \Sigma, l, t, \psi_1 \tilde{\mathcal{S}}_{[w_1,w_2]} \psi_2) = \tilde{\bigvee}_{A \in SR_l^{[w_1,w_2]}} (\tilde{\bigwedge}_{l' \in A} \tau(\mathcal{M}, \Sigma, l', t, \psi_1) \tilde{\wedge} \tilde{\bigwedge}_{l'' \in B^+(A)} \tau(\mathcal{M}, \Sigma, l'', t, \psi_2))$$

Note the similarity between the structure of β and τ, with operators that refer to SSTL and TSTL respectively. We want to clarify that SSTL results are provided performing SSTL monitoring with a given confidence level. Therefore, we are not talking about confidence level of TSTL results, but about TSTL results, given the confidence level for the SSTL monitoring. With the current definition of TSTL we can derive more operators. We can obtain the operator $\mathcal{P}_{>p}(\varphi)$ as:

$$\mathcal{P}_{>p}(\varphi) := \mathcal{P}_{<1-p}(\neg\varphi)$$

Moreover, the *everywhere* spatial operator $\tilde{\boxdot}_{[w_1,w_2]}$ can be defined as:

$$\tilde{\boxdot}_{[w_1,w_2]} \psi := \neg \tilde{\diamondsuit}_{[w_1,w_2]} \neg\psi$$

This requires ψ to hold in all the locations reachable from the current one with a total cost between w_1 and w_2. The *eventually* $\tilde{\mathcal{F}}^{[t_1,t_2]}$ and the *globally* $\tilde{\mathcal{G}}^{[t_1,t_2]}$ operators are defined as usual:

$$\tilde{\mathcal{F}}^{[t_1,t_2]} \psi := T \tilde{\mathcal{U}}^{[t_1,t_2]} \psi \qquad \tilde{\mathcal{G}}^{[t_1,t_2]} \psi := \neg \tilde{\mathcal{F}}^{[t_1,t_2]} \tilde{\neg}\psi$$

The *eventually* formula holds if ψ becomes true within t_1 and t_2 time units from the current one, while the *globally* formula requires ψ to be satisfied for each time unit in the relative interval $[t_1, t_2]$. As we already presented, TSTL provides an additional level of analysis for evaluation of spatio-temporal properties of estimated satisfaction probabilities of SSTL properties. Hence, there is a crucial difference between both the analysis and the logical operators used in SSTL and TSTL. For example, the following two TSTL properties ψ_1 and ψ_2:

$$\psi_1 := \mathcal{P}_{<p}(\varphi_1 \wedge \varphi_2) \qquad \psi_2 := \mathcal{P}_{<p}(\varphi_1) \mathbin{\widetilde{\wedge}} \mathcal{P}_{<p}(\varphi_2)$$

are intrinsically different and therefore they can take on different truth values. For example, let us assume that we are working with a disease spread model and we have the following SSTL properties on the number of infected agents I:

$$\varphi_1 := I > 5 \qquad \varphi_2 := I > 10$$

Let assume that, for a given disease probability threshold p:

$$\tau(\mathcal{M}, \Sigma, l, t, \mathcal{P}_{<p}(\varphi_1)) = F \qquad \tau(\mathcal{M}, \Sigma, l, t, \mathcal{P}_{<p}(\varphi_2)) = T$$

This can happen if we choose the value of p between the two estimates, outside their respective intervals. Since $\varphi_1 \wedge \varphi_2 \equiv \varphi_2$ then:

$$\tau(\mathcal{M}, \Sigma, l, t, \mathcal{P}_{<p}(\varphi_1 \wedge \varphi_2)) = \tau(\mathcal{M}, \Sigma, l, t, \mathcal{P}_{<p}(\varphi_2)) = T$$

while:

$$\tau(\mathcal{M}, \Sigma, l, t, \mathcal{P}_{<p}(\varphi_1)) \mathbin{\widetilde{\wedge}} \tau(\mathcal{M}, \Sigma, l, t, \mathcal{P}_{<p}(\varphi_2)) = F$$

Moreover, the first could perhaps be derived empirically from observations, but the second is only expressible with the new logical operator $\widetilde{\wedge}$ and the domain $\mathbb{T}$.

Monitoring the three-valued semantics of the bounded surround

To evaluate the validity of TSTL formulas we implemented monitoring algorithms for each logical operator, structured in a similar way to SSTL monitoring [2]. We illustrate now the monitoring algorithm for the TSTL *bounded surround* operator, which is more elaborate than the other procedures. Given a location $\hat{l}$ and a TSTL bounded surround formula $\psi = \psi_1 \underset{[w_1,w_2]}{\widetilde{\mathcal{S}}} \psi_2$, the algorithm returns the piecewise constant approximation $s_{\psi,\hat{l}}$ of the function that maps each time t with $\tau(\mathcal{M}, \Sigma, \hat{l}, t, \psi_1 \underset{[w_1,w_2]}{\widetilde{\mathcal{S}}} \psi_2)$, in the discrete time set $\mathcal{T}$. The cardinality of this set $\mathcal{T}$ depends on the given SSTL and TSTL formulas; it is the shortest finite sequence of time-steps for which we have the values of the satisfaction function of all the formulas involved[1]. As shown in Algorithm 1, as the first step of the algorithm, we compute the value $s_{\psi_1,l}$ for all the locations $l : 0 \leq w(\hat{l}, l) \leq w_2$ and the value $s_{\psi_2,l}$ for all the locations $l : w_1 \leq w(\hat{l}, l) \leq w_2$. These values are obtained by recursive invocation of the monitoring algorithm on the TSTL subformulas ψ_1 and ψ_2. We set these values for the other locations to be F, $\forall t \in \mathcal{T}$.

[1] We need to take into account that a temporal formula looks T_f time units into the future, hence the domain $[0, T]$ becomes $[0, T - T_f]$.

After this initial step, we iteratively compute a fixed-point function, on the set of locations satisfying the cost bounds, to get the value of the bounded surround formula, for each time step in the discrete time set $\mathcal{T}$. This fixed-point coincides with the limit of the sequence $(\chi_i)_{i\in\mathbb{N}}$, $\chi_i : L \to \mathbb{T}$, defined as follows:

1. $\chi_0(l) = s_{\psi_1,l}(t)$
2. $\chi_{i+1}(l) = \chi_i(l) \mathbin{\tilde{\wedge}} (\tilde{\bigwedge}_{l':(l,l')\in E} (\chi_i(l') \mathbin{\tilde{\vee}} s_{\psi_2,l'}(t)))$

where i indicates the iteration. The upper bound on the number of iterations of the algorithm is given by the diameter d_G of the graph; given $\chi(l)$ the fixed point of $\chi_i(l)$, then $\chi(l) = \chi_{d_G+1}(l)$, $\forall l \in L$. The proof of the correctness of the method follows that of the SSTL monitoring. The cost of this computation for each location is $O(d_G|L||\mathcal{T}|)$; therefore, the cost for all locations is $O(d_G|L|^2|\mathcal{T}|)$. For more details, see [2], where a similar approach is used.

Algorithm 1:
Three-Valued Spatio-Temporal Logic: bounded surround operator

```
input: l̂, ψ = ψ1 S[w1,w2] ψ2, 𝒯
forall l ∈ L do
    if 0 ≤ w(l̂, l) ≤ w2 then
        compute s_{ψ1,l};
    else
        if w(l̂, l) ≥ w1 then
            compute s_{ψ2,l};
        else
            s_{ψ2,l} = F
        s_{ψ1,l} = F; s_{ψ2,l} = F
forall t ∈ 𝒯 do
    forall l ∈ L do
        χ_prec(l) = T
        χ(l) = s_{ψ1,l}
    while ∃l ∈ L : χ_prec(l) ≠ χ(l) do
        χ_prec = χ
        forall l ∈ L do
            χ(l) = χ_prec(l) ∧̃ (⋀̃_{l':(l,l')∈E} (χ_prec(l') ∨̃ s_{ψ2,l'}))
    s_{ψ,l̂}(t) = χ(l̂)
return s_{ψ,l̂}
```

4 Modelling and monitoring: MELA and jSSTL

We used the process algebra MELA [14] to formally describe spatial population models and to perform stochastic simulations, in order to produce spatio-temporal trajectories for the SSTL monitoring. This process algebra MELA has

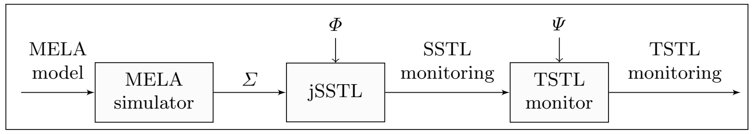

Fig. 1: Σ is the set of spatio-temporal trajectories, Φ the set of SSTL formulas and Ψ the set of TSTL formulas

been developed to build spatial population models of ecological systems, since consideration of the spatial aspect has been recognized as of key importance in ecology. MELA allows one to build models on different discrete spatial structures, to define agent behaviours with spatial constraints on their interactions and probability for these interactions to be effective. Agents can perform different types of actions, that might change their state, their location, or their number in the system. The components in the MELA model generate the states of the underlying stochastic model, a Continuous Time Markov Chain (CTMC) and we perform stochastic simulations using Gillespie's Stochastic Simulation Algorithm (SSA) [15], extracting initial configuration, model structure and parameter values directly from the MELA model description. We chose to use MELA to facilitate the creation of spatial population models since it presents features that fit perfectly with SSTL monitoring settings, such as discrete representation of space and focus on spatial population models. Accordingly, it has been used in order to produce spatio-temporal trajectories, used as input for jSSTL [16], a Java library developed to support monitoring of SSTL properties. Since SSTL works with a discrete space, in particular with weighted graphs, the grid spatial structures in MELA are mapped to a weighted graph structure, to fit with the SSTL framework, with all the weights equal to 1. The results of jSSTL monitoring are used as input to verify TSTL properties. The structure of our spatio-temporal analysis is shown in Figure 1.

5 Case study: Defining safety zones

We now present two case studies, related to fire propagation, using TSTL properties for the identification of safe zones and exit routes. The actual MELA models and more details about the spatio-temporal analysis can be found in `https://ludovicalv.github.io/TSTL`. For the first case study we build a MELA model of forest fire: the spatial structure in this model is a 2D grid, 25×25, with Von Neumann neighbourhood of range 1 and absorbing boundaries. The considered grid is crossed by a road (R), that has a high probability of causing fire (B) in its neighbouring forest. We have zones of particular interest (P) for which we wish to provide strong protection (e.g. picnic areas, houses, regional parks) and zones of safety (S) that will never burn. We want to identify safe areas during the spread of fire. The spread can initiate from the danger zone and it can expand to the neighbouring cells.

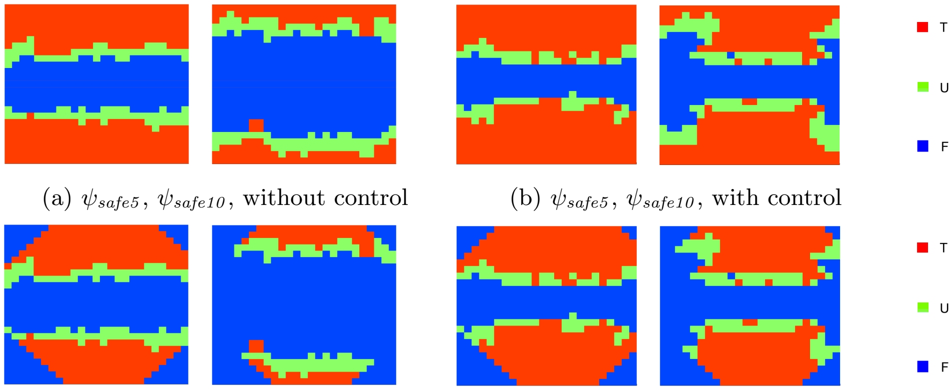

(a) ψ_{safe5}, ψ_{safe10}, without control (b) ψ_{safe5}, ψ_{safe10}, with control

(c) ψ_{area5}, ψ_{area10}, without control (d) ψ_{area5}, ψ_{area10}, with control

Fig. 2: (a, b): validity of TSTL formulas ψ_{safe5}, ψ_{safe10} at $t = 0$ (low risk)
(c, d): validity of TSTL formulas ψ_{area5}, ψ_{area10} at $t = 0$ (safe zones)

- Zones of interest (picnic areas P, brown sign)
- Zones of safety (fire assembly points S, green sign)
- Road (R, black line) and danger zone (D, neighbouring area)
- Green area: vegetation (25 km^2)

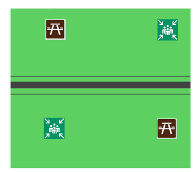

To reach our goal we start with two SSTL properties: the SSTL property φ_{pos}, a "static" property, that predicates about respective position of locations, and the property φ_{fire5}, related to eventually burning (B) over a given time horizon, here for example 5 time steps. With the formula φ_{pos} we identify the locations that connect P and S with a bounded cost, where the cost bounds are chosen given the distance between the two areas, and that are not part of the road. We verify these properties at time 0.

$$\varphi_{pos} := (\Diamond_{[0,17]}(S > 0)) \wedge (\Diamond_{[0,17]}(P > 0)) \wedge (\neg(R > 0))$$

$$\varphi_{fire5} := \mathcal{F}^{[0,5]}(B > 0)$$

We find the safe zones using the TSTL property ψ_{area5}, that identifies the locations satisfying the position requirements and with low probability of burning.

$$\psi_{pos} := \mathcal{P}_{>0.01}(\varphi_{pos}) \qquad \psi_{safe5} := \mathcal{P}_{<0.2}(\varphi_{fire5}) \qquad \psi_{area5} := \psi_{pos} \mathbin{\widetilde{\wedge}} \psi_{safe5}$$

The results of this analysis are shown in Figure 2a and 2c. Using TSTL we can also identify the zone of higher risk: we use the operator *everywhere* to identify the locations for which all the closest neighbours have high probability of being on fire, within the first 10 time steps (ψ_{risk}).

$$\varphi_{fire10} := \mathcal{F}^{[0,10]}(B > 0) \qquad \psi_{fire10} := \mathcal{P}_{>0.8}(\varphi_{fire10}) \qquad \psi_{risk} := \underset{[1,1]}{\widetilde{\boxdot}}\, \psi_{fire10}$$

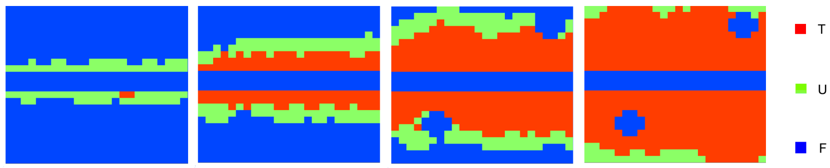

(a) $t = 1$ (b) $t = 5$ (c) $t = 10$ (d) $t = 15$

Fig. 3: TSTL property ψ_{risk}: high risk zones, for $t = 1, 5, 10, 15$ time steps

The results are shown in Figure 3. We can observe that the locations neighbouring the road and the safe zones do not satisfy ψ_{risk}: in fact, not all their neighbouring locations will burn, since both the road and the safe zones will never catch fire. We can also introduce a control measure, like a firebreak [17], to protect the areas between the zone of interest and the safe zones. Firebreaks work as a barrier and they usually consist of a gap in vegetation that slows down the fire spreads. For this reason, we add in our model fire detectors that, once the fire is detected at a given distance, will activate the control measure. These control actions will reduce the probability of fire spread, e.g. cutting the vegetation in the neighbouring area. Using MELA we can also estimate the expense (accumulated performance cost) associated with the control measures, such as the cumulative reward for Continuous Stochastic Reward Logic [18], and use these estimations to balance between effectiveness and expense of the control actions. The difference between the two models, without/with control, is shown in Figure 2, where we can observe a wider safe zone in the model with the control measure. We verify the TSTL properties ψ_{pos} (for position) and ψ_{safe5}, ψ_{safe10} (for low risk of fire spread in 5 and 10 time steps, respectively), while ψ_{area5} and ψ_{area10} are used to define the safe zones in the different scenarios:

$$\psi_{safe5} := \mathcal{P}_{<0.2}(\varphi_{fire5}) \qquad \psi_{safe10} := \mathcal{P}_{<0.2}(\varphi_{fire10})$$

$$\psi_{area5} := \psi_{pos} \mathbin{\widetilde{\wedge}} \psi_{safe5} \qquad \psi_{area10} := \psi_{pos} \mathbin{\widetilde{\wedge}} \psi_{safe10}$$

We can expand our analysis using TSTL to identify the limited areas of risk, in situations where we have isolated danger spots spread in the area, using the TSTL *bounded surround* operator and the formula:

$$\psi_{riskArea} := (\mathcal{P}_{>0.8}(B > 0)) \, \widetilde{\mathcal{S}}_{[w_1,w_2]} (\mathcal{P}_{<0.2}(B > 0))$$

We can also analyse the effectiveness of the control measures, verifying that the risk probabilities do not exceed a given threshold over time.

$$\psi_{lowRisk} := \widetilde{\mathcal{G}}^{[0,t]} \mathcal{P}_{<0.2}(B > 0)$$

In both case studies, for each property we run 30 simulations to perform TSTL verification. With these relatively few runs we have an overall insight about the dynamics and about the differences between distinct models, with a simple representation of complex systems and properties. Nevertheless, this analysis takes into account the uncertainty, that can be used to determine the need of more runs for more precise results. We will expand our current framework to automatically provide additional simulations to refine the analysis of the atomic propositions, when a more precise result is needed.

6 Case study: Emergency evacuation route

For the second case study we introduce another fire related model: we want to identify the most appropriate fire exit in a situation where the evacuation routes are already defined. We aim to identify the safe evacuation routes from the centre of the grid to the assembly points, located in the corners, having a fire starting in the location in the lower left corner, that can spread to the neighbouring locations. We build a MELA model on a 2D grid (25×25) where the fire can spread everywhere, apart from the assembly points (safe zones). In the parts of the grid defined as a route we have two agents in parallel, one that identifies the presence of people and the other one that identifies the presence of fire.

- Fire spread model (on a grid, 25 km^2)
- Agents at the center
- Different exit routes (grey lines) to safe zones (located in the corners)
- For each route: ($fire$ || $people$)

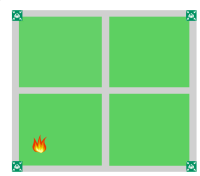

In this example the movement of people and fire are modelled separately, they do not influence each other in the model. We will gather both types of information in the study of TSTL properties. This model focusses more on the concrete safe path and on actual movement, while in the previous one we were defining the different safe areas in time, without specifying the route. In the inflammable area the $fire$ agent can be on fire (B, burning) or not (I, inflammable) while the exit route locations can be empty (EM), occupied (Occ) or passed (P); P represents a cell that was occupied but empty again. To identify the safe evacuation routes, we use TSTL to identify cells that have low probability of being on fire (ψ_{fire}) and non-zero probability of being occupied (ψ_{occ}), given the agent movement in the model.

$$\varphi_{occ} := Occ > 0 \qquad \varphi_{fire} := B > 0$$

$$\psi_{occ} := \mathcal{P}_{>0.01}(\varphi_{occ}) \qquad \psi_{fire} := \mathcal{P}_{<0.2}(\varphi_{fire}) \qquad \psi_{safe} := \psi_{occ} \,\tilde{\wedge}\, \psi_{fire}$$

The verification output of TSTL property ψ_{safe} shows the routes that will lead to the assembly point safely, as shown in Figure 4. To be able to identify the safe evacuation routes from the beginning, instead of observing their temporal evolution, we can check the TSTL property $\psi_{safeRoute}$ at $t = 0$. We want to identify the route that, with probability higher than 0.8, will not be on fire if occupied, in this case in the first 10 time-steps.

$$\varphi_{route} := (EM > 0) \vee (Occ > 0) \vee (P > 0)$$

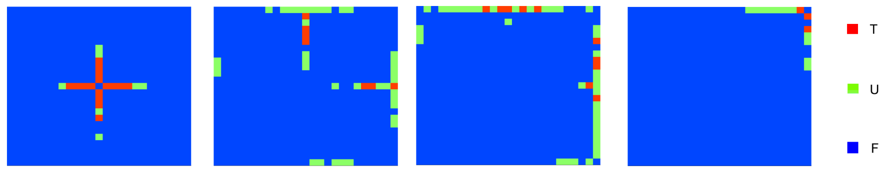

(a) $t = 1$ (b) $t = 5$ (c) $t = 7.5$ (d) $t = 10$

Fig. 4: Temporal evolution of safe evacuation routes: TSTL property ψ_{safe}, with movement rate of the agents equal to 2.0

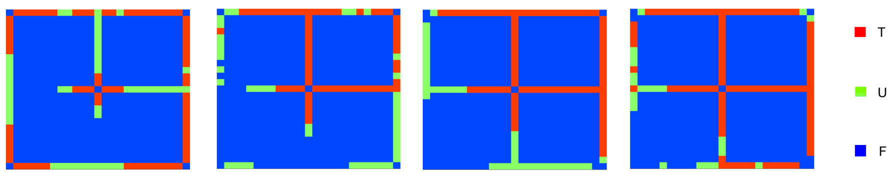

(a) $r = 1.0$ (b) $r = 2.0$ (c) $r = 3.0$ (d) $r = 4.0$

Fig. 5: Safe evacuation routes: TSTL property $\psi_{safeRoute}$, with different movement rate of the agents

$$\varphi_{notFire} := \neg((Occ > 0) \wedge (B > 0)) \equiv \neg(Occ > 0) \vee \neg(B > 0)$$
$$\varphi_{Gsafe} := \mathcal{G}^{[0,10]}(\varphi_{route} \wedge \varphi_{notFire}) \qquad \psi_{safeRoute} := \mathcal{P}_{>0.8}(\varphi_{GSafe})$$

We will check this TSTL property changing the rate of agent movement in the MELA model, as shown in Figure 5. We can observe that if the rate of movement is not high enough, there are not safe options to reach the assembly points. As further analysis, we examine the number of unknown values over time, given TSTL properties and changing the quantity of spatio-temporal trajectories to analyse. We study the percentage of locations having unknown as truth value for different formulas, in both case studies, for different numbers of simulation runs. We observed that the percentage decreases with the increase of the number of runs. Since the width of the confidence intervals depends to a large extent on this value, with an increase in the number of runs we tend to give a more precise estimation of the satisfaction probability. Therefore we have narrower confidence intervals as input for TSTL monitoring and a consequent smaller percentage of unknown values. The three-valued approach is useful to discriminate among TSTL properties in the process of acquiring spatio-temporal trajectories, until the satisfaction set is large enough.

7 Discussion

In the current framework development we use verification of SSTL formulas as input for TSTL monitoring and the starting point for the spatio-temporal analysis. We want to point out that TSTL can be used to predicate on estimated satisfaction probabilities of formulas specified with other logics and also on more general uncertain values with an estimated confidence, as long as the required format is maintained (estimated value for each location at each time point).

As future case studies we will apply our framework to model the spread of invasive species, in particular giant hogweed [19]: we will analyse the effectiveness of different control measures to protect areas of interest, such as regional parks, taking into account also the suitability of the different locations for plant colonisation. In particular, we will analyse the difference between prevention (control outside the boundaries of the area) and direct action (eradication when the invasive species are detected inside the area), considering also the expense associated with the different measures.

As a future extension for TSTL we will define and implement the operator *bounded reachable* $\widetilde{\mathcal{R}}$. This operator can be seen as a spatial until with direction and associated with a path. We will be able to verify properties related with locations reachable within a given cost range and satisfying defined TSTL properties, and the existence of a connecting path formed only by locations satisfying a given set of TSTL properties. In the case studies we presented, the use of this new operator would have allowed us to identify safe paths without having to mimic the actual movement, detecting different possible solutions. Using this new TSTL operator we could verify if *there is a safe location (assembly point S) that we can reach passing only through locations with low probability of burning, with a cost w, $w \in [w_1, w_2]$*:

$$\psi := (\mathcal{P}_{<0.2}(B > 0)) \underset{[w_1,w_2]}{\widetilde{\mathcal{R}}} (\mathcal{P}_{>0.01}(S > 0))$$

8 Conclusions

In this paper we presented Three-Valued Spatio-Temporal Logic (TSTL), an extension of Signal Spatio-Temporal Logic (SSTL) that allows us to widen the analysis of spatio-temporal properties of stochastic systems. We have shown how this extension is used to study the spatio-temporal evolution of the estimated satisfaction probabilities of given SSTL formulas. We implemented the monitoring algorithms for each TSTL operator and used them in the case studies to perform the novel analysis, checking the validity of different TSTL formulas. We used TSTL to identify the zones that have high risk of catching fire during a fire spread and to find the safest evacuation routes, checking the ones that have high probability to be safe over time. We provide the novel spatio-temporal logic with a three-valued semantics to handle the intrinsic uncertainty related to the statistical methods used to estimate the satisfaction probabilities. The three-valued approach allows us to perform initial analysis with a relatively small set of spatio-temporal trajectories, taking into account the uncertainty; on the other hand, it also provides a decision tool on the number of simulations needed for drawing stronger conclusions.

Acknowledgement

This work was supported by Microsoft Research Cambridge through its PhD Scholarship Programme and by the EU project QUANTICOL 600708. Glenn Marion was funded by the Scottish Government Rural and Environment Science and Analytical Services Division (RESAS).

References

1. A. Legay, B. Delahaye, S. Bensalem, *Statistical Model Checking: An Overview*, pp. 122–135. Springer Berlin Heidelberg, 2010.
2. L. Nenzi, L. Bortolussi, V. Ciancia, M. Loreti, M. Massink, "Qualitative and quantitative monitoring of spatio-temporal properties," in *Runtime Verification - 6th International Conference, RV*, pp. 21–37, 2015.
3. M. Aiello, I. Pratt-Hartmann, J. Van Benthem, *Handbook of Spatial Logics*. Springer Netherlands, 2007.
4. L. Cardelli, A. D. Gordon, "Anytime, anywhere: modal logics for mobile ambients," in *POPL 2000*, pp. 365–377, ACM, 2000.
5. J. Reif, A. Sistla, "A multiprocess network logic with temporal and spatial modalities," *Journal of Computer and System Sciences*, vol. 30(1), pp. 41 – 53, 1985.
6. H. Hansson, B. Jonsson, "A logic for reasoning about time and reliability," *Formal Aspects of Computing*, vol. 6(5), pp. 512–535, 1994.
7. C. Baier, B. Haverkort, H. Hermanns, J. P. Katoen, "Model-checking algorithms for continuous-time markov chains," *IEEE TSE*, vol. 29, pp. 524–541, 2003.
8. S. Gottwald, "Many-valued logic," in *Stanford Encyclopedia of Philosophy*, 2008.
9. J. Łukasiewicz, *Selected Works*. Amsterdam: North-Holland Pub. Co., 1970.
10. S. C. Kleene, "On notation for ordinal numbers," *JSL*, vol. 3(4), pp. 150–155, 1938.
11. J.-P. Katoen, D. Klink, M. Leucker, V. Wolf, "Three-valued abstraction for probabilistic systems," *JLAP*, vol. 81(4), pp. 356 – 389, 2012.
12. K. Sen, M. Viswanathan, G. Agha, *Model-Checking Markov Chains in the Presence of Uncertainties*, pp. 394–410. Springer Berlin Heidelberg, 2006.
13. O. Maler, D. Nickovic, *Monitoring Temporal Properties of Continuous Signals*, pp. 152–166. Springer Berlin Heidelberg, 2004.
14. L. Luisa Vissat, J. Hillston, G. Marion, M. J. Smith, "Mela: Modelling in ecology with location attributes," in *EPTCS*, vol. 227, pp. 82–97, 2016.
15. D. T. Gillespie, "Exact stochastic simulation of coupled chemical reactions," *J. Phys. Chem.*, vol. 81(25), pp. 2340–2361, 1977.
16. L. Nenzi, L. Bortolussi, M. Loreti, "jsstl - a tool to monitor spatio-temporal properties," *ValueTools*, 2016.
17. D. Cerotti, M. Gribaudo, A. Bobbio, C. T. Calafate, P. Manzoni, *A Markovian Agent Model for Fire Propagation in Outdoor Environments*, pp. 131–146. Springer Berlin Heidelberg, 2010.
18. C. Baier, B. Haverkort, H. Hermanns, J.-P. Katoen, "On the logical characterisation of performability properties," in *Automata, Languages and Programming*, vol. LNCS 1853, pp. 780–792, Springer-Verlag, 2000.
19. S. Catterall, A. R. Cook, G. Marion, A. Butler, P. E. Hulme, "Accounting for uncertainty in colonisation times: a novel approach to modelling the spatio-temporal dynamics of alien invasions using distribution data," *Ecography*, vol. 35(10), pp. 901–911, 2012.