

QUALITATIVE AND QUANTITATIVE MONITORING OF SPATIO-TEMPORAL PROPERTIES WITH SSTL *

LAURA NENZI^a, LUCA BORTOLUSSI^b, VINCENZO CIANCIA^c, MICHELE LORETI^d,
AND MIEKE MASSINK^e

^a Institute of Computer Engineering, University of Technology, Vienna, Austria

^b DMG, University of Trieste, Trieste, Italy

^{c,e} Istituto di Scienza e Tecnologie dell'Informazione "A. Faedo" - CNR, Pisa, Italy

^d Scuola di Scienze e Tecnologie, University of Camerino, Camerino, Italy

ABSTRACT. In spatially located, large scale systems, time and space dynamics interact and drives the behaviour. Examples of such systems can be found in many smart city applications and Cyber-Physical Systems. In this paper we present the *Signal Spatio-Temporal Logic* (SSTL), a modal logic that can be used to specify spatio-temporal properties of *linear time* and *discrete space* models. The logic is equipped with a *Boolean* and a *quantitative semantics* for which efficient monitoring algorithms have been developed. As such, it is suitable for real-time verification of both white box and black box complex systems. These algorithms can also be combined with stochastic model checking routines. SSTL combines the until temporal modality with two spatial modalities, one expressing that something is true *somewhere* nearby and the other capturing the notion of being *surrounded* by a region that satisfies a given spatio-temporal property. The monitoring algorithms are implemented in an open source Java tool. We illustrate the use of SSTL analysing the formation of patterns in a Turing Reaction-Diffusion system and spatio-temporal aspects of a large bike-sharing system.

Key words and phrases: Spatio-Temporal Logic, Signal Temporal Logic, Quantitative Semantics, Monitoring Algorithms, Real-Time Verification, Turing Pattern, Bike Sharing, Statistical Model Checking.

* Work partially funded by the EU-FET project QUANTICOL (nr. 600708) and the FRA-UNITS. We thank Diego Latella and Ezio Bartocci for the discussions, EB for sharing the code to generate traces of the Turing Pattern model and Cheng Feng for the data of the Bike Sharing system.



1. INTRODUCTION

The ongoing digital revolution is concretising itself in scenarios where a large number of computational devices, located in space, interact among each other and with humans in an open and mutating environment. Designing and controlling such systems, and guaranteeing some sort of correctness at run-time, are incredibly challenging problems.

In this paper, we consider spatially located systems, where time and space dynamics interact and drive the system behaviour. Examples are Cyber-Physical Systems, like pacemaker devices controlling the rhythm of heart beat, and Collective Adaptive Systems, such as bike sharing systems in smart cities or the guidance of crowd movement in emergency situations.

Controlling and designing spatio-temporal behaviours requires appropriate formal tools to describe such properties, and to monitor and verify whether, and how robustly, they are satisfied by a system. Formal methods play a central role in these tasks, providing both formal languages to specify spatio-temporal models and properties, and algorithms to verify such properties on models and on traces of monitored systems. In this paper we consider large and complex systems for which standard model checking procedures, i.e. those that check in an exhaustive way which states of a system model satisfy a given property, are not feasible. For these kinds of systems simulation and testing are currently the preferred verification methods. This is the area of monitoring (run-time verification) [22, 33], where an individual simulation trace \mathbf{x} of a system is checked against a modal logic formula, using an automatic verification procedure.

Related work. Logical specification and monitoring of *temporal* properties of systems is a well-developed area. Here we mention Signal Temporal Logic (STL) [22, 33], an extension of Metric Interval Temporal Logic (MITL) [3], describing linear-time properties of real-valued signals. STL has monitoring routines both for its Boolean and quantitative semantics. The latter provides a measure of the satisfaction degree of a formula [22, 23, 33]. In monitoring algorithms verification is performed over single trajectories.

Much work can be found also in the area of spatial logic, mostly to reason about continuous space (so-called *topological* spatial logics, see [1]) and focussing on theoretical issues such as expressivity and decidability. Instead, model checking and monitoring procedures have a more recent history. Some logics have been proposed to specify properties of networks of processes [38], for (graph) rewriting theories [4], [35], bigraphs [21] and data structures such as graphs [13] and heaps [12]. However these logics have often high computational cost, and sometimes they are even undecidable.

In the present work, we will focus on a notion of *discrete* space. The reason is that many applications, such as bike sharing systems or meta-population epidemic models [34], are naturally framed in a discrete spatial structure. Moreover, in many circumstances continuous space is abstracted as a grid or as a mesh. This is the case, for instance, in many numerical methods to simulate the spatio-temporal dynamics of Partial Differential Equations (PDE). Hence, this class of models is naturally dealt with by checking properties on such a discretisation.

One of the recent spatial logics that has inspired the work in the current paper is the *Spatial Logic for Closure Spaces* (SLCS) [14, 18]. SLCS is a logic proposed for both a discrete and topological notion of space and is based on the theory of (quasi discrete) closure spaces [27, 28]. In SLCS some interesting spatial operators have been introduced such as

the (unbounded) surround operator and the near-operator inspired by the closure operator of closure spaces and many derived operators among which well-known operators such as *somewhere* and *everywhere*.

Even more challenging is the combination of spatial and temporal operators [1] and few works exist with a practical perspective. Among those, SLCS has been extended with a branching time temporal logic in STLCS [15, 17] leading to an implementation of a spatio-temporal model checker. STLCS has been applied in the context of smart public transportation, such as public buses [16] and bike sharing systems [19, 20]. The temporal aspects are introduced as “snapshot” models, where each snapshot represents the situation of the discrete space at a particular point in time. Differently from the logic considered in this paper, STLCS is only equipped with a Boolean semantics.

SpaTeL [30] is a *linear-time* spatio-temporal logic that combines the temporal logic STL with the *Tree Spatial Superposition Logic* (TSSL) [29]. It has been provided with both a Boolean semantics as well as with a quantitative semantics. The spatial logic TSSL is suited to specify properties of quad trees. These are spatial data structures that are constructed by recursively partitioning a finite space into uniform quadrants. TSSL provides a statistical way to describe the distribution of discrete states in a particular partition of the space. Its spatial logic operators can be used to zoom in and zoom out of particular areas of interest. In this way, very complex spatial structures may be captured, but at the price of a complex formulation of spatial properties. The properties can in practice only be learned from some template image using supervised learning algorithms and are not human-readable, even too long to be displayed. In contrast, the sub-language for spatial properties that we adopt in this paper is descriptive and topological in nature.

Contributions. In this work we present a formal definition of the *Signal Spatio-Temporal Logic* (SSTL). SSTL integrates the temporal modalities of STL with two spatial modalities; the *somewhere* operator and a novel *bounded* version of the topological *surrounded* operator. The surrounded operator is inspired by the topological spatial until operator of SLCS [14], adding useful metric bounds on spatial distances. A key point of our logic is that the spatial properties are not only considered as atomic propositions of a temporal logic as in SpaTeL. This means that the temporal and spatial operators can be arbitrarily nested, permitting the specification of complex spatio-temporal behaviours. The logic is provided with a Boolean and a quantitative semantics, both equipped with efficient monitoring algorithms. The major challenge is to monitor the bounded surrounded operator for the quantitative semantics, for which we propose a novel fixed point algorithm, discussing in detail its correctness and computational cost. A prototype tool has been developed in **Java** and can be downloaded from <https://github.com/Quanticol/jsstl>.

This article is an extension of the conference papers [8, 36]. Specifically, the description of the monitoring algorithms has been enhanced, describing the procedure in detail, including proofs of all statements (in particular, correctness of the monitoring algorithm for the *surrounded* operator). The Turing Reaction-Diffusion case study of [36] has been used as a running example throughout the paper, to illustrate the intuition and formal semantics of the new spatial operators. We show how SSTL can deal with stochastic scenarios, considering the effects of external perturbations on the Turing system, adding a white Gaussian noise to the set of equations. To this purpose SSTL monitoring algorithms are combined with statistical model checking techniques, following up on earlier works on such combinations with STL [6]. Furthermore, we illustrate the application of SSTL, and the verification

methods we have developed, on a larger case study concerning issues of the spatio-temporal distribution of bikes in a bike-sharing system, modelled as a Continuous Time Markov Chain (CTMC).

Paper structure. The paper is organised as follows: Section 2 introduces some background concepts on STL and on discrete topologies. Section 3 presents the syntax and the semantics of SSTL and the Turing system running example. Section 4 introduces the monitoring algorithms and the implementation details of the model checker. Section 5 describes the extension of the monitoring procedure to stochastic systems and its applications on the running example. In Section 6, the logic SSTL is applied to a Bike-Sharing system, while conclusions are drawn in Section 7. Proofs of the main results are provided in Appendix A.

2. BACKGROUND MATERIAL

In this section we provide some general background material and notation directly relevant to the results developed in subsequent sections.

2.1. Weighted undirected graphs. We consider discrete models of space that can be represented as finite undirected graphs. Edges of such graphs are equipped with a positive weight, providing a metric structure to the space in terms of shortest path distances. The weight will often represent the distance between two nodes. This is the case, for instance, when the graph is a discretisation of continuous space. However, the notion of weight is more general, and may be used to encode different kinds of information. As an example, in a model where nodes represent locations in a city and edges represent streets, the weight could represent the average travelling time, which can be different between two paths with the same physical length but different levels of congestion or different number of traffic lights.

Definition 2.1. A (positive) **weighted undirected graph** is a tuple $G = (L, E, w)$, where:

- L is a finite set of locations (nodes), $L \neq \emptyset$
- $E \subseteq L \times L$ is a symmetric relation, namely the set of connections between nodes (edges),
- $w : E \rightarrow \mathbb{R}_{>0}$ is a function that returns the positive cost/weight of each edge.

We will use both $(\ell, \ell') \in E$ and $\ell E \ell'$ as equivalent notations for an edge in the relation E . The space is equipped with a distance metric.

Definition 2.2. For $\ell, \ell' \in L$, the **weighted distance** is defined as

$$d(\ell, \ell') := \min\left\{\sum_{e \in \sigma} w(e) \mid \sigma \text{ is a path between } \ell \text{ and } \ell'\right\}.$$

This means that the weighted distance is a metric that returns the cost of the shortest path, for each pair of nodes of the graph; where the shortest path is the path that minimises the sum of the weights of the edges of the path.

Remark 2.3. Let E^* be the set containing all the pairs of connected locations, i.e. the transitive closure of E . If L is finite, and if we define an order on the locations, $L = \{\ell_1, \dots, \ell_i, \dots\}$, then the weighted distance can be seen as a matrix $(d)_{(\ell_i, \ell_j) \in E^*}$, where $d[i, j]$ is the distance between ℓ_i and ℓ_j .

Furthermore, we denote by $L_{[d_1, d_2]}^\ell$ the set of locations ℓ' at a distance between d_1 and d_2 from ℓ , formally

$$L_{[d_1, d_2]}^\ell := \{\ell' \in L \mid d_1 \leq d(\ell, \ell') \leq d_2, \text{ with } d_1, d_2 \geq 0\}.$$

2.2. Closure spaces. In this work we focus on finite graphs as an algorithmically tractable representation of space. However, *spatial* logics traditionally use more abstract structures, very often of a topological nature (see [1] for an exhaustive reference). Indeed, the logic we propose can also be defined in a more abstract setting. We can frame a generalised notion of topology on graphs within the so-called *Cech closure spaces*, a superclass of topological spaces allowing a clear formalisation of the semantics of the spatial surrounded operator on both topological and graph-like structures (see [14, 15, 18] and the references therein). As an example, we mention the topological notion of *external boundary* of a set of nodes A , instantiated on weighted graphs as the set of nodes directly connected to an element of A but not part of it.

Definition 2.4. Given a subset of locations $A \subseteq L$, we define the *external boundary* of A as:

$$B^+(A) := \{\ell \in L \mid \ell \notin A \wedge \exists \ell' \in A \text{ s.t. } (\ell', \ell) \in E\}.$$

2.3. Signal Temporal Logic. *Signal Temporal Logic* (STL) [22, 33] is a linear dense time-bounded temporal logic that extends *Metric Interval Temporal Logic* (MITL) [3] with a set of atomic propositions $\{\mu_1, \dots, \mu_m\}$ that specify properties of real valued traces, therefore mapping real valued traces into Boolean values.

Let $\mathbf{x} : \mathbb{T} \rightarrow \mathbb{D}$ be a trace that describes an evolution of our system, where $\mathbb{T} = \mathbb{R}_{\geq 0}$ represents continuous time and $\mathbb{D} = \mathbb{D}_1 \times \dots \times \mathbb{D}_n \subseteq \mathbb{R}^n$ is the domain of evaluation; then each $\mu_j : \mathbb{D} \rightarrow \mathbb{B}$ is of the form $\mu_j(x_1, \dots, x_n) \equiv (f_j(x_1, \dots, x_n) \geq 0)$, where $f_j : \mathbb{D} \rightarrow \mathbb{R}$ is a (possibly non-linear) real-valued function and $\mathbb{B} = \{true, false\}$ is the set of Boolean values. The projections $x_i : \mathbb{T} \rightarrow \mathbb{D}_i$ on the i^{th} coordinate/variable are called the *primary signals* and, for all j , the function $s_j : \mathbb{T} \rightarrow \mathbb{R}$, defined by point-wise application of f_j to the image of \mathbf{x} , namely $s_j(t) := f_j(x_1(t), \dots, x_n(t))$, is called the *secondary signal* [23].

The syntax of STL is given by

$$\varphi := \mu \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \mathcal{U}_{[t_1, t_2]} \varphi_2,$$

where μ is an atomic proposition, conjunction and negation are the standard Boolean connectives, $[t_1, t_2]$ is a real positive dense interval with $t_1 < t_2$ and $\mathcal{U}_{[t_1, t_2]}$ is the *bounded until* operator. The latter specifies that formula φ_1 holds until, at time $t \in [t_1, t_2]$, formula φ_2 is satisfied. This operator can be used to define the operators *eventually* and *always*. The *eventually* operator $\mathcal{F}_{[t_1, t_2]}$ and the *always* operator $\mathcal{G}_{[t_1, t_2]}$ can be defined as usual with \top denoting *true*: $\mathcal{F}_{[t_1, t_2]}\varphi := \top \mathcal{U}_{[t_1, t_2]}\varphi$, $\mathcal{G}_{[t_1, t_2]}\varphi := \neg \mathcal{F}_{[t_1, t_2]}\neg\varphi$. Formula $\mathcal{F}_{[t_1, t_2]}\varphi$ states that φ is eventually satisfied at a time $t \in [t_1, t_2]$, while $\mathcal{G}_{[t_1, t_2]}\varphi$ indicates that at each time $t \in [t_1, t_2]$, φ is satisfied.

3. SSSL: SIGNAL SPATIO-TEMPORAL LOGIC

Signal Spatio-Temporal Logic (SSSL) is a spatial extension of STL [22, 33] with two spatial modalities: the *bounded somewhere* operator $\diamond_{[d_1, d_2]}$ and the *bounded surrounded* operator $\mathcal{S}_{[d_1, d_2]}$. SSSL is interpreted on spatio-temporal, real-valued signals. In this section, we first introduce the signals and then present the syntax and the Boolean and quantitative semantics of SSSL, using a pattern formation model as running example.

3.1. Spatio-Temporal Signals. We define signals with continuous time and discrete space. In particular, the space is represented by a weighted undirected graph $G = (L, E, w)$, defined in Section 2, while the time domain \mathbb{T} will usually be the real-valued interval $[0, T]$, for some $T > 0$. Formally, a spatio-temporal signal, is a function $\mathbf{s} : \mathbb{T} \times L \rightarrow \mathbb{D}$, where L is the set of locations and \mathbb{D} is the domain of evaluation. \mathbb{D} is a subset of $\mathbb{R}^* = \mathbb{R} \cup \{+\infty, -\infty\}$. Signals with $\mathbb{D} = \mathbb{B} = \{0, 1\}$ are called Boolean signals, whereas those where $\mathbb{D} = \mathbb{R}^*$ are called real-valued or quantitative signals.

A spatio-temporal trace is a function $\mathbf{x} : \mathbb{T} \times L \rightarrow \mathbb{R}^n$ s.t. $\mathbf{x}(t, \ell) = (x_1(t, \ell), \dots, x_n(t, \ell)) \in \mathbb{D}$, where each $x_i : \mathbb{T} \times L \rightarrow \mathbb{R}$, for $i = 1, \dots, n$, is the projection on the i^{th} coordinate/variable. Note that these projections have the form of quantitative signals. They are called the *primary signals* of the trace. We can thus see the trace as a set of primary signals. This means that SSSL can be used to specify spatio-temporal properties of such traces. Spatio-temporal traces can be obtained by simulating a stochastic model or by computing the solution of a deterministic system. In previous work [8] the framework of patch-based population models is discussed, which generalise population models and are a natural setting from which both stochastic and deterministic spatio-temporal traces of the considered type emerge. An alternative source of traces are measurements of real systems. For the purpose of this work, it is irrelevant which is the source of traces, as we are interested in their off-line monitoring.

Spatio-temporal traces are then converted into spatio-temporal Boolean or quantitative signals in the following way: similarly to the case of STL, each *atomic predicate* μ_j is of the form $\mu_j(x_1, \dots, x_n) \equiv (f_j(x_1, \dots, x_n) \geq 0)$, for some function $f_j : \mathbb{D} \rightarrow \mathbb{R}$, where (x_1, \dots, x_n) are the primary signals. Each atomic proposition gives rise to a spatio-temporal signal. In the Boolean case, one may define function $s_j : \mathbb{T} \times L \rightarrow \mathbb{B}$; given a trace \mathbf{x} , this gives rise to the Boolean signal $s_j(t, \ell) = \mu_j(\mathbf{x}(t, \ell))$ by point-wise lifting. Similarly, a quantitative signal is obtained as the real-valued function $s_j : \mathbb{T} \times L \rightarrow \mathbb{R}$, with $s_j(t, \ell) = f_j(\mathbf{x}(t, \ell))$. These signals, derived from the atomic proposition, are called the *secondary signals*.

When the space L is finite, as in our case, we can represent a spatio-temporal signal as a finite collection of temporal signals. More specifically, the signal $s(t, \ell)$ can be equivalently represented by the collection $\{s_\ell(t) \mid \ell \in L\}$. We will stick mostly to this second notation in the following, as it simplifies the presentation.

3.2. SSSL Syntax. The syntax of SSSL is given by

$$\varphi := \mu \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \mathcal{U}_{[t_1, t_2]} \varphi_2 \mid \diamond_{[d_1, d_2]} \varphi \mid \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2.$$

Atomic predicates, Boolean operators, and the time bounded until operator $\mathcal{U}_{[t_1, t_2]}$ are those of STL. The new spatial operators are the *somewhere* operator, $\diamond_{[d_1, d_2]}$, and the *bounded surrounded* operator $\mathcal{S}_{[d_1, d_2]}$, where $[d_1, d_2]$ is a closed real interval with $d_1 < d_2$. We can derive also the *everywhere* operator as $\boxplus_{[d_1, d_2]} \varphi := \neg \diamond_{[d_1, d_2]} \neg \varphi$.

The somewhere and the everywhere operators were inspired by the modal operators of the *Multiprocess Network Logic* [38]; the idea originated from the necessity to describe behaviours at a certain distance from a specific point, e.g., “from a bike sharing station, in a radius of 100 meters, there are more than 30 bikes” or “in all the positions around my location, at a distance less than 1 km, there are no infected individuals”. Formally, the spatial somewhere operator $\diamond_{[d_1, d_2]}\varphi$ requires φ to hold in a location reachable from the current one with a cost greater than or equal to d_1 and less than or equal to d_2 . The cost is given by the shortest weighted distance between the locations, i.e. the sum of the weights of the edges of the shortest path (Def. 2.2). We use the word “cost” to distinguish it from the classical spatial notion of distance. Indeed, we can have two streets with the same distance but different travel time due to the presence of traffic lights or congestion. In Figure 1, representing a regular grid where the distance between adjacent nodes is 1, we provide some examples of spatial properties. In the graph of the figure, the orange point satisfies the property $\diamond_{[3,5]}$ *pink*. Indeed, there exists a point at a distance 4 from the orange point that satisfies the pink property. The *everywhere* operator $\boxminus_{[d_1, d_2]}\varphi := \neg \diamond_{[d_1, d_2]} \neg \varphi$ requires φ to hold in *all* the locations reachable from the current one with a total cost between d_1 and d_2 . In Figure 1, the orange point of the graph satisfies the property $\boxminus_{[2,3]}$ *yellow*. Indeed, all the points at a distance between 2 and 3 from the orange point satisfy the yellow property.

The (bounded) surrounded operator $\varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2$ is inspired by the *surrounded* operator of the *Spatial Logic for Closure Spaces* SLCS (see [14, 18]), and it is a spatial interpretation of the temporal *until* connective. It expresses the topological notion of being surrounded by a φ_2 -region, while being in a φ_1 -region, with additional metric constraints. Informally speaking, the intended meaning of $\varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2$ is that one cannot escape from a φ_1 -region without passing from a node that satisfies φ_2 and, in any case, one has to reach a φ_2 -node at a distance between d_1 and d_2 . An example drawn from analysis of bike sharing facilities in *smart cities* is the property “there are no bikes in the station where I am, but all the bike stations directly connected with this one have at least one bike available, and are located at a distance less than 100 meters from here”.

The surrounded operator makes the logic strictly more expressive. Whenever its distance bounds are trivial (i.e., equal to $[0, \infty)$), it allows us to express *global* properties, depending upon the state of the whole system and not only on the local neighbourhood of each point. This version of the surrounded operator with non-trivial metric bounds adds an additional level of expressivity, allowing us to restrict the attention to subregions of the space defined by the distance constraints. In any case, the topological properties captured by the surrounded operator cannot be captured by the simpler somewhere and everywhere operators, which in turn cannot be defined in terms of the surrounded operator, due to the effect of metric bounds.

Formally, the bounded variant of the surrounded formula, i.e. $\varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2$, is true in a location ℓ , when ℓ belongs to a set of locations A satisfying φ_1 and at a distance less than d_2 from ℓ , the external boundary $B^+(A)$ of A must contain only locations satisfying φ_2 . Furthermore, locations in $B^+(A)$ must be reached from ℓ with a cost between d_1 and d_2 . $B^+(A)$ is the set of all the locations that do not belong to A but that are directly connected with a location in A . In Figure 1, the green points satisfy *green* $\mathcal{S}_{[0, 100]}$ *blue*. Indeed, for each green point we can find a region that contains the point, such that all its points are green and all the points connected with an element that belongs to the region are blue and satisfy the metric constraint. Instead, the property *green* $\mathcal{S}_{[2, 3]}$ *blue* is satisfied only by

the dark green point. The reason is that such a dark green point is the only point for which there exists a region (the green region) such that all its elements are at a distance less than 3 from it and are green; and all the elements of the external boundary (the blue region) are at a distance between 2 and 3 from it.




Figure 1: Example of spatial properties. The orange point satisfies $\diamond_{[3,5]}$ *pink*. The orange point satisfies $\square_{[2,3]}$ *yellow*. All green points satisfy *green* $\mathcal{S}_{[0,100]}$ *blue*. The dark green point satisfies also *green* $\mathcal{S}_{[2,3]}$ *blue*.

3.3. Running Example: Turing Patterns. To illustrate more complex examples of spatial and spatio-temporal properties in SSTL we introduce a model of pattern formation in a reaction-diffusion system which will be used as a running example in the following sections.

Alan Turing conjectured in [39] that pattern formation is a consequence of the coupling of reaction and diffusion phenomena involving different chemical species, and that these can be described by a set of PDE reaction-diffusion equations, one for each species. The natural analogue, systems of agents interacting and moving in continuous space, is however prohibitively expensive to analyse computationally; an approach that is more amenable to analysis is to discretise space into a number of cells which are assumed to be spatially homogeneous, and to replace spatial diffusion with transitions between different cells.

From the point of view of formal verification, the formation of patterns is an inherently spatio-temporal phenomenon, in that the relevant issue is how the spatial organisation of the system changes over time.

Pattern formation model. Our model, similar to that in [29,30], describes the production of skin pigments that generate spots in animal furs. The reaction-diffusion system is discretised, according to a Finite Difference scheme, as a system of ODEs whose variables are organised in a $K \times K$ rectangular grid. More precisely, we treat the grid as a weighted undirected graph, where each cell $(i, j) \in L = \{1, \dots, K\} \times \{1, \dots, K\}$ is a location (node), edges connect each pairs of neighbouring nodes along four directions (so that each node has at most 4 adjacent nodes), and the weight of each edge is always equal to the spatial length-scale δ of the system¹. We consider two species (chemical substances) A and B in a

¹For simplicity, here we fix $\delta = 1$. Note that using a non-uniform mesh the weights of the edges of the resulting graph will not be uniform.

$K \times K$ regular grid, obtaining the system:

$$\begin{cases} \frac{dx_{i,j}^A}{dt} = R_1 x_{i,j}^A x_{i,j}^B - x_{i,j}^A + R_2 + D_1(\mu_{i,j}^A - x_{i,j}^A) & i = 1.., K, j = 1, \dots, K, \\ \frac{dx_{i,j}^B}{dt} = R_3 x_{i,j}^A x_{i,j}^B + R_4 + D_2(\mu_{i,j}^B - x_{i,j}^B) & i = 1.., K, j = 1, \dots, K, \end{cases} \quad (3.1)$$

where: $x_{i,j}^A$ and $x_{i,j}^B$ are the concentrations of the two species in the cell (i, j) ; $R_i, i = 1, \dots, 4$ are the parameters that define the reaction between the two species; D_1 and D_2 are the diffusion constants; $\mu_{i,j}^A$ and $\mu_{i,j}^B$ are the inputs for the (i, j) cell, that is

$$\mu_{i,j}^n = \frac{1}{|\nu_{i,j}|} \sum_{\nu \in \nu_{i,j}} x_{\nu}^n \quad n \in \{A, B\}, \quad (3.2)$$

where $\nu_{i,j}$ is the set of indices of cells adjacent to (i, j) . The spatio-temporal trace of the system is the function $\mathbf{x} = (x^A, x^B) : [0, T] \times L \rightarrow \mathbb{R}^{K \times K} \times \mathbb{R}^{K \times K}$ where each x^A and x^B are the projection on the first and second variable, respectively. In Figure 2 the concentration of species A is shown for a number of time points, generated by the numerical integration of System 3.1. The initial conditions have been set randomly to concentrations of both species in a range between values 0 and 16. It can be observed that at times $t = 20$ and $t = 50$ the shape of the pattern appears and then remains stable. Clearly, some regions (in blue) have a low concentration of A surrounded by regions with a high concentration of A. We consider the regions with low concentration of species A as the ‘spots’ in the pattern. The opposite happens for the value of B (high density regions surrounded by low density regions, not shown).

3.4. SSSL Boolean Semantics. We first define the Boolean semantics for SSSL. This semantics, as customary, returns true/false depending on whether the observed trace satisfies the SSSL specification.

Definition 3.1 (SSSL Boolean semantics). The Boolean satisfaction relation for an SSSL formula φ over a spatio-temporal trace \mathbf{x} is given by:

$$\begin{aligned} (\mathbf{x}, t, \ell) \models \mu & \Leftrightarrow \mu(\mathbf{x}(t, \ell)) = 1 \\ (\mathbf{x}, t, \ell) \models \neg\varphi & \Leftrightarrow (\mathbf{x}, t, \ell) \not\models \varphi \\ (\mathbf{x}, t, \ell) \models \varphi_1 \wedge \varphi_2 & \Leftrightarrow (\mathbf{x}, t, \ell) \models \varphi_1 \wedge (\mathbf{x}, t, \ell) \models \varphi_2 \\ (\mathbf{x}, t, \ell) \models \varphi_1 \mathcal{U}_{[t_1, t_2]} \varphi_2 & \Leftrightarrow \exists t' \in t + [t_1, t_2] : (\mathbf{x}, t', \ell) \models \varphi_2 \wedge \forall t'' \in [t, t'], (\mathbf{x}, t'', \ell) \models \varphi_1 \\ (\mathbf{x}, t, \ell) \models \diamond_{[d_1, d_2]} \varphi & \Leftrightarrow \exists \ell' \in L : (\ell', \ell) \in E^* \wedge d_1 \leq d(\ell', \ell) \leq d_2 \wedge (\mathbf{x}, t, \ell') \models \varphi \\ (\mathbf{x}, t, \ell) \models \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2 & \Leftrightarrow \exists A \subseteq L_{[0, d_2]}^\ell : \ell \in A \wedge \forall \ell' \in A, (\mathbf{x}, t, \ell') \models \varphi_1 \\ & \quad \wedge B^+(A) \subseteq L_{[d_1, d_2]}^\ell \wedge \forall \ell'' \in B^+(A), (\mathbf{x}, t, \ell'') \models \varphi_2. \end{aligned}$$

A trace \mathbf{x} satisfies φ in location ℓ , denoted by $(\mathbf{x}, \ell) \models \varphi$, if and only if $(\mathbf{x}, 0, \ell) \models \varphi$.

Example 3.2 (Spot Formation properties). We illustrate the Boolean semantics of SSSL and, in particular, the use of the surrounded operator to characterise spot and pattern formation of the running example introduced in Section 3.3. In order to classify spots, sub-regions of the grid have to be identified that present a high (or low) concentration of one of the species, surrounded by a low (high, respectively) concentration of the same species. For example, one can capture the spots of the A species using the spatial formula

$$\varphi_{\text{spot}} := (x^A \leq h) \mathcal{S}_{[d_1, d_2]} (x^A > h). \quad (3.3)$$




Figure 2: Value of x^A for the system (3.1) for $t = 0, 5, 7, 12, 20, 50$ time units with parameters $K = 32, R_1 = 1, R_2 = -12, R_3 = -1, R_4 = 16, D_1 = 5.6$ and $D_2 = 25.5$. The initial condition has been set randomly in a range between values 0 and 16. The colour map for the concentration is specified in the legend on the right of each subplot.

A trace \mathbf{x} satisfies φ_{spot} at time t , in the location (i, j) , $(\mathbf{x}, t, (i, j)) \models \varphi_{\text{spot}}$, if and only if there is a subset $L' \subset L$, that contains (i, j) , such that all elements in L' have a distance less than d_2 from (i, j) , and x^A , at time t , is less or equal to a constant value h . Furthermore, in each element in the boundary of L' we have that $x^A > h$ at time t , and its distance from (i, j) is in the interval $[d_1, d_2]$. Note that the use of distance bounds in the surrounded operator allows one to constrain the size (diameter) of the spot. Recall that we are considering a spatio-temporal system, so this spatial property alone is not enough to describe the formation of a pattern *over time*; to identify the insurgence time of the pattern and whether it remains stable over time we combine the spatial property with temporal operators in the following way:

$$\varphi_{\text{spotFormation}} := \mathcal{F}_{[T_{\text{pattern}}, T_{\text{pattern}} + \delta]} \mathcal{G}_{[0, T_{\text{end}}]}(\varphi_{\text{spot}}); \quad (3.4)$$

$\varphi_{\text{spotFormation}}$ states that eventually, at a time between T_{pattern} and $T_{\text{pattern}} + \delta$, the property φ_{spot} becomes true and remains true for at least T_{end} time units. Figure 3(b) shows the satisfaction of the property $\varphi_{\text{spotFormation}}$ in each cell $(i, j) \in L$ for a simulated trace, for the Boolean semantics with formula parameters set to $d_1 = 1$ and $d_2 = 6$. In Figure 3(d) we similarly show the satisfaction of the same property for a smaller value of the maximal dimension of a spot $d_2 = 4$ (the other parameters are as indicated in the caption of Figure 3). The plot shows the satisfaction at time $t = 0$ because by default we have that $(\mathbf{x}, \ell) \models \varphi$, if and only if $(\mathbf{x}, 0, \ell) \models \varphi$. It is clear that the satisfaction of the property is capable of identifying which locations belong to the spots and which not. If we make the distance

constraint stricter, by reducing the width of the interval $[d_1, d_2]$, we are able to identify only the “centre” of the spot, as shown in Figure 3 (d). However, in this case we may fail to identify spots that have an irregular shape (i.e., that deviate too much from a circular shape).

Example 3.3 (Pattern Formation properties). The next formula, φ_{pattern} , describes the global spatial pattern in which every location is part of a spot or has a nearby spot. This is expressed by the following SSTL formula:

$$\varphi_{\text{pattern}} := \boxplus_{[0, d_{\text{max}}]} \blacklozenge_{[0, d_{\text{spot}}]} \varphi_{\text{spot}}, \quad (3.5)$$

where \blacklozenge and \boxplus are the everywhere and somewhere operators, d_{max} is chosen to cover the whole space, and d_{spot} specifies the maximal distance between spots. Checking this formula in a random location of our space is enough to verify the presence of the pattern; this is so because the first part of the formula, $\boxplus_{[0, d_{\text{max}}]}$, covers all the locations of the grid. This is an example of how one can describe a global property also with a semantics that verifies properties in single locations. For $d_{\text{max}} = 45$ and $d_{\text{spot}} = 15$ (and the other parameters are as indicated in Figure 3) for a solution of the system (3.1) we obtain the result *true* for the Boolean semantics. Changing the diffusion constants D_1 and D_2 affects the shape and size of the spots or disrupts them, as we can see in Figure 4. We evaluate the pattern formula (3.5) with parameters as in Figure 3, for the patterns in Figure 4(a) and (b), where $D = [1.5, 23.6]$ and $D = [8.5, 40.7]$, respectively, and the other parameters equal to the previous model. This gives the result *false* in both cases. Formula (3.4), though, is still true in some locations. This is due to the irregularity of the spots (where, as in Figure 4(a), some spots can have a shape similar to the model in Figure 3 (a)), or due to particular boundary effects on the border of the grid (where fractions of spots still remain, as in Figure 4(a)).

3.5. Quantitative Semantics. The quantitative semantics returns a *real value* that can be interpreted as a measure of the strength with which the specification is satisfied or falsified by an observed trajectory. More specifically, the sign of such a satisfaction score is related to the truth of the formula (positive stands for true), while the absolute value of the score is a measure of the robustness of the satisfaction or dissatisfaction. This definition of quantitative measure is based on [22, 23], and it is a reformulation of the robustness degree of [25].

Definition 3.4 (SSTL Quantitative Semantics). The quantitative satisfaction function $\rho(\varphi, \mathbf{x}, t, \ell)$ for an SSTL formula φ over a spatio-temporal trace \mathbf{x} is given by:

$$\begin{aligned} \rho(\mu, \mathbf{x}, t, \ell) &= f(\mathbf{x}(t, \ell)) \quad \text{where } \mu \equiv (f \geq 0) \\ \rho(\neg\varphi, \mathbf{x}, t, \ell) &= -\rho(\varphi, \mathbf{x}, t, \ell) \\ \rho(\varphi_1 \wedge \varphi_2, \mathbf{x}, t, \ell) &= \min(\rho(\varphi_1, \mathbf{x}, t, \ell), \rho(\varphi_2, \mathbf{x}, t, \ell)) \\ \rho(\varphi_1 \mathcal{U}_{[t_1, t_2]} \varphi_2, \mathbf{x}, t, \ell) &= \sup_{t' \in t + [t_1, t_2]} (\min(\rho(\varphi_2, \mathbf{x}, t', \ell), \inf_{t'' \in [t, t']} (\rho(\varphi_1, \mathbf{x}, t'', \ell)))) \\ \rho(\blacklozenge_{[d_1, d_2]} \varphi, \mathbf{x}, t, \ell) &= \max\{\rho(\varphi, \mathbf{x}, t, \ell') \mid \ell' \in L, (\ell', \ell) \in E^*, d_1 \leq d(\ell', \ell) \leq d_2\} \\ \rho(\varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2, \mathbf{x}, t, \ell) &= \max_{A \subseteq L_{[0, d_2]}^\ell, \ell \in A, B^+(A) \subseteq L_{[d_1, d_2]}^\ell} (\min(\min_{\ell' \in A} \rho(\varphi_1, \mathbf{x}, t, \ell'), \min_{\ell'' \in B^+(A)} \rho(\varphi_2, \mathbf{x}, t, \ell''))) \end{aligned}$$

where ρ is the quantitative satisfaction function, returning a real number $\rho(\varphi, \mathbf{x}, t)$ quantifying the degree of satisfaction of the property φ by the trace \mathbf{x} at time t . Moreover, $\rho(\varphi, \mathbf{x}, \ell) := \rho(\varphi, \mathbf{x}, 0, \ell)$.




Figure 3: Satisfaction of formula (3.4) with $h = 0.5$, $T_{\text{pattern}} = 19$, $\delta = 1$, $T_{\text{end}} = 30$, $d_1 = 1$, $d_2 = 6$ for (b), (c) and $d_2 = 4$ for (d). (a) Concentration of A at time $t = 50$; (b) (d) Boolean semantics of the property $\varphi_{\text{spotFormation}}$; the cells (locations) that satisfy the formula are in red, the others are in blue; (c) Quantitative semantics of the property $\varphi_{\text{spotFormation}}$; The value of the robustness is given by a colour map as specified in the legend on the right of the figure.




Figure 4: Snapshots at time $t = 50$ of x^A for the model (3.1) with parameters with parameters $K = 32$, $R_1 = 1$, $R_2 = -12$, $R_3 = -1$, $R_4 = 16$, and $D = [1.5, 23.6]$ in (a) and $D = [8.5, 40.7]$ in (b).

The definition for the surrounded operator is essentially obtained from the Boolean semantics by replacing conjunctions and universal quantifications with the minimum and disjunctions and existential quantifications with the maximum, as done in [22, 23] for STL.

Remark 3.5 (Soundness Property). As for STL, the quantitative semantics of SSSL is sound with respect to the Boolean semantics. It means the ρ is positive whenever the property holds in the Boolean semantics, and that it has a negative value if the property does not hold. This can be proved by induction on the structure of the formula, whereas it is already shown for STL operators (atomic propositions, Boolean and temporal operators, for fixed locations). The new spatial operators are defined as compositions of maximum and minimum functions (on finite sets and for a fixed time). Therefore, it is easy to prove that the property holds also for them. Indeed, considering for example $\min\{\rho(\varphi, \mathbf{x}, t, \ell_1), \rho(\varphi, \mathbf{x}, t, \ell_2)\}$, by induction $\rho(\varphi, \mathbf{x}, t, \ell_1)$ and $\rho(\varphi, \mathbf{x}, t, \ell_2)$ satisfy the soundness property. If they are both positive, they both satisfy φ and also the minimum is positive. If one $\rho(\varphi, \mathbf{x}, t, \ell_i)$ is negative, it does not satisfy φ and the minimum is negative. A similar consideration can be done for the maximum.

Remark 3.6 (Correctness Property). The quantitative semantics satisfies the correctness property with respect to the Boolean semantics. It means that for each formula φ and $\forall \mathbf{x}_1, \mathbf{x}_2$, it holds:

$$(\varphi, \mathbf{x}_1, t, \ell) \models \varphi \text{ and } \|\mathbf{x}_1 - \mathbf{x}_2\|_\infty < \rho(\varphi, \mathbf{x}_1, t, \ell) \Rightarrow (\varphi, \mathbf{x}_2, t, \ell) \models \varphi$$

Recalling that the new spatial operators are defined as compositions of maximum and minimum functions (on finite set and for fixed time), also this property can be proved by induction on the structure of the formula. It holds for atomic predicates, Boolean and temporal operators by the correctness proof of the quantitative semantics for MITL [24]. We need just to prove that it holds for maximum and minimum functions of the robustness function with fixed time. Let $\rho(\psi, \mathbf{x}_1, \ell) = \min\{\rho(\varphi, \mathbf{x}_1, \ell_1), \rho(\varphi, \mathbf{x}_1, \ell_2)\}$ with φ satisfying correctness by induction. If $\|\mathbf{x}_1 - \mathbf{x}_2\|_\infty < \rho(\psi, \mathbf{x}_1, \ell)$ and $\rho(\psi, \mathbf{x}_1, \ell) > 0$ (soundness property), i.e. $(\mathbf{x}_1, \ell_i) \models \varphi$ for $i = 1, 2$, by induction on φ and the minimum function we have that $(\mathbf{x}_2, \ell_i) \models \varphi$ for $i = 1, 2$, then $(\mathbf{x}_2, \ell) \models \psi$; similarly it can be proved for the maximum.

Remark 3.7. The quantitative semantics is a measure of how robust is the truth value of a formula φ for a given trace \mathbf{x} . More specifically, it tells us the maximum size of a perturbation of the secondary signal to preserve the satisfaction of φ [23] (a variant of our definition, [25] returns a robustness measure for the primary signal, i.e. the trace \mathbf{x}). This perturbation can be applied to the quantitative value of the secondary signal at any point in time and space. As such, the value of the robustness defines a tube in the space of secondary signals, uniformly with respect to space-time, containing equisatisfiable trajectories (with respect to the Boolean semantics): any trace in this tube will have the same truth value of φ as \mathbf{x} . We stress that perturbations are applied pointwise in space and time: this version of the quantitative semantics tells us nothing about the effect of time and space warps. This would require a different notion of quantitative semantics, as done e.g. in [2, 23] for time only, coupled with a meaningful notion of perturbation of a discrete space (e.g. on edge weights).

Example 3.8 (Robustness of Spot Formation). As a first example, consider again the property $\varphi_{\text{spotFormation}}$. Figure 3(b) shows the Boolean satisfaction of the property $\varphi_{\text{spotFormation}}$ in each cell $(i, j) \in L$, for $d_1 = 1$ and $d_2 = 6$ and Figure 3(c) shows the result

of the *quantitative* semantics for the same formula, showing for each cell (i, j) the value of the robustness with which the formula is satisfied. Positive values indicate cells where the formula holds, negative values where it does not. Where the value is positive its value amounts to 0.3. This low value is due to the specific choice of the threshold value $h = 0.5$.

Example 3.9 (Robustness of Pattern Formation). The second example considers the global formula φ_{pattern} . This formula resulted in *false* applying the Boolean semantics for both patterns in Figure 4. The quantitative semantics gives -0.05 as a result in both cases, indicating a weak robustness.

Example 3.10 (Perturbation). A strength of spatio-temporal logics is the possibility to nest the temporal and spatial operators. We illustrate this in the following scenario as a variant of the running example. We set as initial conditions for the dynamical system (3.1) its stable state, i.e. the concentrations of substances A and B at time 50 (see Figure 3(a)). We introduce a small perturbation by changing a single value in a specific location in the centre of a spot. The idea is to study the effect of this perturbation over time by checking whether it will disrupt the system or not. Specifically, we perturb the cell $(6, 6)$ in Figure 3(a), by setting $x_{6,6}^A(0) = 10$ while its original value was 0. Dynamically, the perturbation is quickly absorbed and the system returns to the previous steady state. Formally, we consider the following property:

$$\varphi_{\text{pert}} := (x^A \geq h_{\text{pert}}) \wedge (\varphi_{\text{absorb}} \mathcal{S}_{[d_m, d_M]} \varphi_{\text{no.effect}}); \quad (3.6)$$

The meaning of φ_{pert} is that the induced perturbation remains confined inside the original spot when the property is satisfied. More in detail, a trace \mathbf{x} satisfies φ_{pert} in the location (i, j) , i.e. $(\mathbf{x}, (i, j)) \models \varphi_{\text{pert}}$, if and only if $x_{i,j}^A(0) \geq h_{\text{pert}}$, with $h_{\text{pert}} = 10$, (the location is perturbed) and if there is a subset $L' \subseteq L$ that contains location (i, j) such that all its elements have a distance less than d_M from (i, j) and satisfy $\varphi_{\text{absorb}} = \mathcal{F}_{[0, T_p]} \mathcal{G}_{[0, T_d]}(x^A < h')$; φ_{absorb} states that the perturbation of x^A is absorbed within T_p units of time, stabilising back to a value $x^A < h'$ for an additional period of T_d time units. Here h' is a suitable threshold capturing the fact that concentrations have returned close to their value before the perturbation, and it is set to $h' = 3$ in our case. Furthermore, within distance $[d_m, d_M]$ from the original perturbation, where d_M is chosen such that we are within the spot at cell $(6, 6)$ of the non-perturbed system, $\varphi_{\text{no.effect}} := \mathcal{G}_{[0, T]}(x^A < h')$ is satisfied; i.e., no relevant effect is observed, the value of x^A stably remains below h' . In Figure 5, we report the evaluation of the quantitative semantics for φ_{pert} , zooming in on the 15×15 lower left corner of the original grid. As shown in the figure, the perturbed location $(6, 6)$ satisfies the property.

It is interesting to observe that this property requires a genuine nesting of space and time modalities, and as such it cannot be expressed in logics like SpaTeL, in which the spatial properties are limited to the level of atomic propositions.

4. MONITORING ALGORITHMS

In this section, we present the offline monitoring algorithms to check the satisfaction of a formula φ on a trace $\mathbf{x}(t, \ell)$. The monitoring procedures spatially extend the property monitors introduced in [33] for the Boolean and in [22] for the quantitative semantics of STL.




Figure 5: Boolean (a) and quantitative (b) semantics for formula 3.6 with parameters $h_{\text{pert}} = 10$, $d_m = 1$, $d_M = 2$, $T_p = 1$, $T_d = 10$, $h' = 3$, and $T = 20$.

As for STL formulae, our algorithms work with a bottom-up approach on the syntax tree of φ , iteratively computing the temporal signals of each subformula. Each node of the tree represents a subformula, the leafs are the atomic propositions and the root represents the complete formula. Given a spatio-temporal trace $\mathbf{x}(t, \ell)$, the algorithm starts computing the spatio-temporal Boolean/quantitative signals of all the atomic propositions, then it moves upwards in the tree computing the spatio-temporal Boolean/quantitative signals of a node using the signals of its child and a specific algorithm for each operator of the logic. Finally, the spatial Boolean/quantitative satisfaction function corresponds to the value of the signal at time zero $\rho_\varphi(0, \ell)$. An example of the procedure is shown in Figure 6.

In the case of the Boolean semantics, for each subformula ψ , the algorithm constructs a spatio-temporal signal s_ψ s.t. $s_\psi(\ell, t) = 1$ iff the subformula is true in position ℓ at time t . In the case of the quantitative semantics, for each subformula ψ , the signal s_ψ corresponds to the value of the quantitative satisfaction function ρ , for any time t and location ℓ . Here, we discuss in detail the algorithms to check the new spatial operators: the somewhere and surrounded operators; the procedures for the other Boolean and temporal operators are similar to STL and will be just briefly recalled.

The processing of the somewhere operator is a simple extension of the disjunction operator. The treatment of the bounded surrounded modality $\psi = \varphi_1 \mathcal{S}_{[w_1, w_2]} \varphi_2$, instead, deviates substantially from all these procedures and will be discussed in more detail. In particular, in the following, we will present two recursive algorithms to compute the Boolean and the quantitative satisfaction, assuming the Boolean/quantitative signals of φ_1 and φ_2 being known.

We remark that this surround operator requires very different algorithms from those developed for timed modalities, as space is bi-directional, thus it makes sense to consider both *reaching* and *being reached*; classical path-based model checking does not coincide with spatial model checking also because loops in space are not relevant in the definition of *surrounded* operators.

4.1. Boolean Monitoring. The algorithm proceeds inductively bottom-up on the parse tree of the formula. Given a formula φ , to determine if $(\mathbf{x}, \ell) \models \varphi$, we construct for all ψ that are subformulas of φ , a Boolean signal $s_\psi : [0, T] \times L \rightarrow \mathbb{B}$ s.t. $s_\psi(t, \ell) = 1$ iff $(\mathbf{x}, t, \ell) \models \psi$ and 0 otherwise. At the termination of the algorithm, we obtain the signal $s_\varphi(t, \ell)$ whose value at $t = 0$ determines whether the trace \mathbf{x} satisfies φ in location ℓ (at time 0). The properties

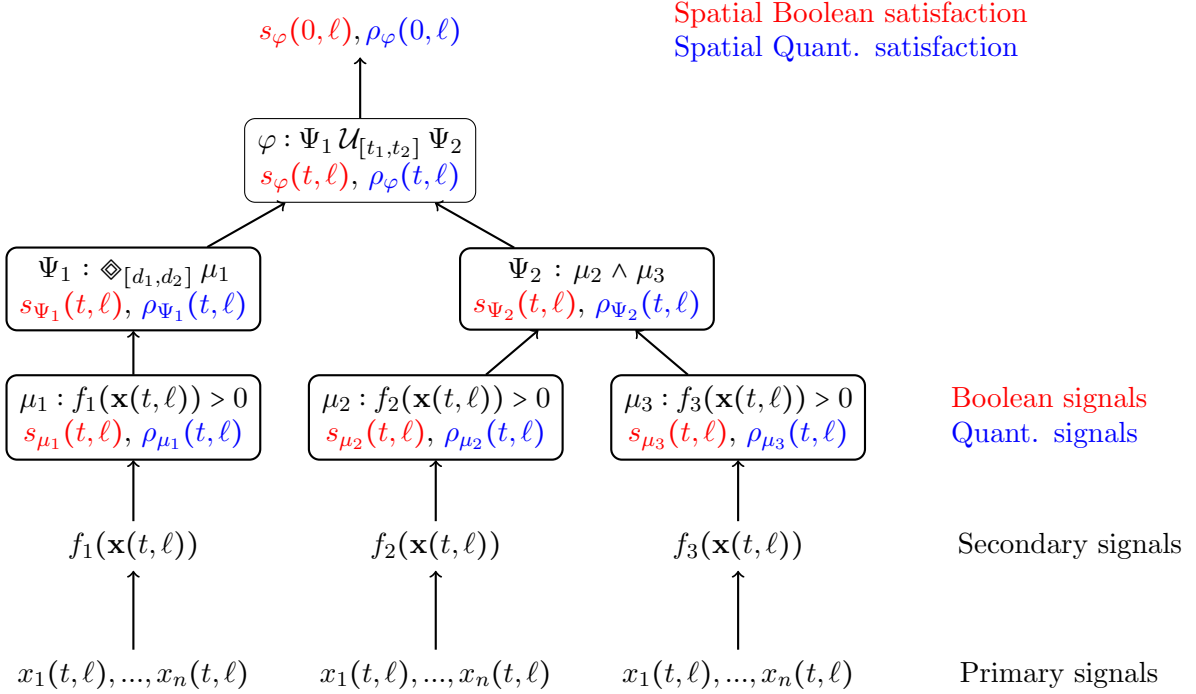


Figure 6: The monitoring procedure of the SSTL formula $\varphi : (\diamond_{[d_1, d_2]} \mu_1) \mathcal{U}_{[t_1, t_2]} (\mu_2 \wedge \mu_3)$.

can be verified pointwise for each location and each time independently, indeed $(\mathbf{x}, \ell, t) \models \varphi$ means “the trace \mathbf{x} in location ℓ at time t satisfies the property φ ”.

To optimise the monitoring procedure, we split the time domain according to the *minimal interval covering* $\mathcal{I}_{s_1, \dots, s_n}$ consistent with a set of temporal Boolean signals s_1, \dots, s_n (as in [33]) that we describe below. A temporal Boolean signal is a function $s : [0, T] \rightarrow \mathbb{B}$. Note that, we can represent the signal $s_\psi : [0, T] \times L \rightarrow \mathbb{B}$ as a finite collection of temporal signals $\{s_{\psi, \ell}\}_{\ell \in L}$ where $s_{\psi, \ell}(t) := s_\psi(\ell, t)$.

Definition 4.1. Given a time interval I , and a set of temporal signals s_1, \dots, s_n with $s_i : I \rightarrow \mathbb{B}$, the *minimal interval covering* $\mathcal{I}_{s_1, \dots, s_n}$ of I consistent with the set of signals s_1, \dots, s_n is the shortest finite sequence of left-closed right-open intervals I_1, \dots, I_h such that $\bigcup_j I_j = I$, $I_i \cap I_j = \emptyset$, $\forall i \neq j$, and for $k \in \{1, \dots, n\}$, $s_k(t) = s_k(t')$ for all t, t' belonging to the same interval². Restricting to a single signal s , the *positive minimal interval covering* of s is $\mathcal{I}_s^+ = \{I \in \mathcal{I}_s \mid \forall t \in I : s(t) = 1\}$. The *negative minimal interval covering* of s is $\mathcal{I}_s^- = \{I \in \mathcal{I}_s \mid \forall t \in I : s(t) = 0\}$, and it holds $\mathcal{I}_s = \mathcal{I}_s^+ \cup \mathcal{I}_s^-$ and $\mathcal{I}_s^+ \cap \mathcal{I}_s^- = \emptyset$.

The positive interval covering $\mathcal{I}_{s_{\psi, \ell}}^+$ corresponds to the *satisfaction set* of the formula over the signal $s_{\psi, \ell}$. Furthermore, any signal can be written as $s = s_1 \vee s_2 \vee \dots \vee s_k$ where each s_i is a *unitary signal*, meaning that it has a singleton positive interval, i.e., $\mathcal{I}_{s_i}^+ = \{[t_1, t_2)\}$ for some $t_1 < t_2 \in \mathbb{R}_{\geq 0}$. Then, we have that $\mathcal{I}_s = \mathcal{I}_{s_1, \dots, s_k}$ and $\mathcal{I}_s^+ = \bigcup_i \mathcal{I}_{s_i}^+$. The idea is that the satisfaction set of a formula over a signal can be seen as the union of disjoint intervals.

²The fact that we can always obtain a finite interval covering is a consequence of the restriction to closed intervals $[t_1, t_2]$, $t_1 < t_2$, in STL. Further details about signals and interval covering are provided in [33].

Using these definitions of signals, interval coverings, and satisfaction set, the procedure for the classic operators of STL is similar to the one described in [33]. We briefly recall these procedures in the following and then we describe the algorithms for the new spatial operators.

Atomic Predicates. $\psi = \mu$. The computation of the Boolean signal associated with an atomic predicate is a direct application of Definition 3.1: $s_{\mu,\ell}(t) = \mu(\mathbf{x}(t, \ell))$.

Negation. $\psi = \neg\varphi$, then $\mathcal{I}_{s_{\neg\varphi,\ell}}^+ = \mathcal{I}_{s_{\varphi,\ell}}^-$.

Disjunction. $\psi = \varphi_1 \vee \varphi_2$, then, given $s_{\varphi_1,\ell}, s_{\varphi_2,\ell}$, let \mathcal{I} be the minimal interval covering consistent with *both* signals. For each $I_i \in \mathcal{I}$, we construct the signal $s_{\psi,\ell}(I_i) = s_{\varphi_1,\ell}(I_i) \vee s_{\varphi_2,\ell}(I_i)$ and we merge adjacent positive intervals to obtain $\mathcal{I}_{\psi,\ell}^+$.

Until. $\psi = \varphi_1 \mathcal{U}_{[a,b]} \varphi_2$. As we are working with future temporal modalities, we need to shift intervals *backwards*. This has to be done independently for each unitary signal, then taking the union of the so obtained satisfaction sets. Given two unitary signals p and q , the signal $\psi = p \mathcal{U}_{[a,b]} q$ is the unitary signal such that $\mathcal{I}_{\psi}^+ = \{((I_p \cap I_q) \ominus [a,b]) \cap I_p\}$, where $[m,n] \ominus [a,b] = [m-b, n-a] \cap [0, T]$ is the Minkowski sum. In the general case, let $s_{\varphi_1,\ell} = p_1 \vee \dots \vee p_n$ and $s_{\varphi_2,\ell} = q_1 \vee \dots \vee q_m$ be signals written as union of unitary signals, then $\psi = s_{\varphi_1,\ell} \mathcal{U}_{[a,b]} s_{\varphi_2,\ell} = \bigvee_{i=1}^n \bigvee_{j=1}^m p_i \mathcal{U}_{[a,b]} q_j$. The proof of this result can be found in [33].

Somewhere. $\psi = \diamond_{[d_1,d_2]} \varphi$. As remarked at the beginning of this section, and relying on the fact that we have a finite number of locations, we can process each location in the signal independently. Given the signal $s_{\psi,\ell}$, for a *fixed location* ℓ , we can rewrite the somewhere operator as a disjunction between all signals in locations ℓ' s.t. $d_1 \leq d(\ell', \ell) \leq d_2$. This allows us to use the monitoring procedure for disjunction, constructing the minimal interval covering \mathcal{I} consistent with all $s_{\varphi,\ell'}$ signals s.t. $d_1 \leq d(\ell', \ell) \leq d_2$, and defining, for each $I_i \in \mathcal{I}$:

$$s_{\psi,\ell}(I_i) = \bigvee_{d_1 \leq d(\ell', \ell) \leq d_2} s_{\varphi,\ell'}(I_i).$$

The satisfaction set $\mathcal{I}_{s_{\psi,\ell}}^+$ is then the union of the positive I_i (i.e., I_i s.t. $s_{\psi,\ell}(I_i) = 1$), merging adjacent positive intervals.

We stress here that the rewriting of the somewhere operator as a finite disjunction is possible only at monitoring time, when the space structure is known. In fact, one cannot express the somewhere operator in terms of the or-operator in general, as this requires the knowledge of the spatial model to be verified, both in terms of locations and of distances among them. Therefore, the spatial somewhere operator is not merely syntactic sugar. Additionally, it can be applied so to countably infinite discrete spaces, and it can be generalised to continuous spaces [18]. Furthermore, even if we assume that the space is finite and if we encode the space structure of the model in the formula, expanding the operator as a disjunction would produce a blowup of the size of the formula which is *exponential* in the nesting level of spatial operators, and hence it would result in an exponential increase in the complexity of the monitoring procedure.

Algorithm 1 Boolean monitoring for the surrounded operator

```

1: input  $\ell, \psi = \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2$ 
2:  $\forall \ell' \in L_{[0, d_2]}^\ell$  compute  $s_{\varphi_1, \ell'}, s_{\varphi_2, \ell'}$ 
3: compute  $\mathcal{I}_{s_{\psi, \ell}}$  {the minimal interval covering consistent with  $s_{\varphi_1, \ell'}, s_{\varphi_2, \ell'}, \ell' \in L_{[0, d_2]}^\ell$ }
4: for all  $I_i \in \mathcal{I}_{s_{\psi, \ell}}$  do
5:    $V = \{\ell' \in L_{[0, d_2]}^\ell \mid s_{\varphi_1, \ell'}(I_i) = 1\}$ 
6:    $Q = \{\ell' \in L_{[d_1, d_2]}^\ell \mid s_{\varphi_2, \ell'}(I_i) = 1\}$ 
7:    $W = B^+(Q \cup V)$ 
8:   while  $W \neq \emptyset$  do
9:      $W' = \emptyset$ 
10:    for all  $\ell \in W$  do
11:       $N = \text{pre}(\ell) \cap V = \{\ell' \in V \mid \ell E \ell'\}$ 
12:       $V = V \setminus N$ 
13:       $W' = W' \cup (N \setminus Q)$ 
14:    end for
15:     $W = W'$ 
16:  end while
17:   $s_{\psi, \ell}(I_i) = \begin{cases} 1 & \text{if } \ell \in V, \\ 0 & \text{otherwise.} \end{cases}$ 
18: end for
19: merge adjacent positive intervals  $I_i$ , i.e.,  $I_i$  s.t.  $s_{\psi, \ell}(I_i) = 1$ 
20: return  $s_{\psi, \ell}$ 

```

Surrounded. $\psi = \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2$. Algorithm 1 presents the procedure to monitor the Boolean semantics of ψ at location ℓ , returning the temporal Boolean signal $s_{\psi, \ell}$ of ψ at location ℓ . The algorithm first computes the set of locations $L_{[0, d_2]}^\ell$ that are at distance d_2 or less from ℓ , and then, recursively, the temporal Boolean signals $s_{\varphi_1, \ell'}$ and $s_{\varphi_2, \ell'}$, for $\ell' \in L_{[0, d_2]}^\ell$. These signals provide the satisfaction of the sub-formulas φ_1 and φ_2 at each point in time, and for each location of interest. Then, a minimal interval covering consistent with all the signals $s_{\varphi_1, \ell'}$ and $s_{\varphi_2, \ell'}$ is computed, and to each such interval, a core procedure similar to that of [14] is applied. More specifically, we first compute the set of locations W in which both φ_1 and φ_2 are false, and that belong to the external boundary of the set of locations that satisfy φ_1 (V) or φ_2 (Q). The locations in W are “bad” locations, that cannot be part of the external boundary of the set A of φ_1 -locations which has to be surrounded exclusively by φ_2 -locations. Hence, the main loop of the algorithm removes iteratively from V all those locations that have a neighbour in W (set N , line 13), constructing a new set T containing only those locations in N that do not satisfy φ_2 , until a fixed point is reached. As each location can be added to W and is processed only once, the complexity of the algorithm is linear in the number of locations and linear in the size of the interval covering. A correctness theorem, stated below, can be proven in a similar way as in [14]. The proof is reported in Appendix A.1.

Theorem 4.2. *Given a graph $G = (L, w, E)$, two properties φ_1 and φ_2 , a trace \mathbf{x} and a location ℓ , let $s_{\psi, \ell} = \text{BoolSurround}(G, \mathbf{x}, \varphi_1, \varphi_2, \ell)$ and $\mathcal{I}_{s_{\psi, \ell}}$ be the minimal interval covering consistent with $\{s_{\varphi_1, \ell'}, s_{\varphi_2, \ell'}\}_{\ell' \in L_{[0, w_2]}^\ell}$, then, for all $I_i \in \mathcal{I}_{s_{\psi, \ell}}$*

$$s_{\psi, \ell}(I_i) = 1 \iff (\mathbf{x}, t, \ell) \models \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2 \quad \forall t \in I_i$$

4.2. Quantitative Monitoring. We now turn to the monitoring algorithm for the quantitative semantics. As the input signals are functions of continuous time, we need to make some assumptions to represent them in a finite way. Contrary to other approaches, like [23], which assume signals to be piecewise linear, we make a simpler assumption, namely that they are piecewise constant functions. More specifically, we discretise time with a step h , and consider a piecewise constant representation of a signal \mathbf{x} , assuming its value between time $t_k = kh$ and $t_{k+1} = (k+1)h$ being equal to $\mathbf{x}(kh)$. We will show how to compute the quantitative semantics for such an approximation, and discuss the error we introduce in case the input signal is Lipschitz continuous.

Monitoring Boolean operators is straightforward, we just need to apply the definition of the quantitative semantics pointwise in the discretisation, i.e. at each time step kh . The time bounded until operator can also be easily computed by replacing the min and max over dense real intervals in its definition by the corresponding min and max over the corresponding finite grid of time points. In this case, however, an error is introduced, due to the discrete approximation of the Lipschitz continuous signal in intermediate points. The error accumulates at a rate proportional to Mh , where M is the Lipschitz constant of the signal \mathbf{x} , formally defined in Proposition 4.3.

Monitoring the somewhere operator $\diamond_{[d_1, d_2]} \varphi$ is also immediate: once the location ℓ of interest is fixed, similarly to the Boolean semantics, we can just turn it into a disjunction of the signals $s_{\varphi, \ell'}$ for each location $\ell' \in L_{[d_1, d_2]}^\ell$.

The only non-trivial monitoring algorithm is the one for the spatial surrounded operator, which we discuss below. However, as the satisfaction score is computed at each time point of the discretisation and depends on the values of the signals at that time point only, this algorithm introduces no further error w.r.t. the time discretisation. Hence, we can globally bound the error introduced by the time discretisation (see the Appendix for the proof):

Proposition 4.3. *Let the primary signal \mathbf{x} be Lipschitz continuous, as well as the functions defining the atomic predicates. Let M be a Lipschitz constant for all secondary signals, and h be the discretisation step. Given a SCTL formula φ , let $u(\varphi)$ count the number of temporal until operators in φ , and denote by $\rho(\varphi, \mathbf{x})$ its satisfaction score over the trace \mathbf{x} and by $\rho(\varphi, \hat{\mathbf{x}})$ the satisfaction score over the discretised version $\hat{\mathbf{x}}$ of \mathbf{x} with time step h . Then $\|\rho(\varphi, \mathbf{x}) - \rho(\varphi, \hat{\mathbf{x}})\| \leq u(\varphi)Mh$.*

The quantitative monitoring procedure for the bounded surrounded operator is shown in Algorithm 2. Similarly to the Boolean case, the algorithm for the surrounded formula $\psi = \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2$ takes as input a location ℓ and returns the quantitative signal $s_{\psi, \ell}$, or better its piecewise constant approximation with time-step h (an additional input, together with the signal duration $T = mh$). As a first step, it computes recursively the quantitative satisfaction signals of subformula φ_1 , for all locations $\ell' \in L_{[0, d_2]}^\ell$, and of subformula φ_2 , for all locations $\ell' \in L_{[d_1, d_2]}^\ell$. Furthermore, it sets all the quantitative signals for φ_1 and φ_2 for the other locations to the constant signal equal to minus infinity. The algorithm runs a fixpoint computation for each time instant in the discrete time set $\{0, h, 2h, \dots, mh\}$. The procedure is based on computing a function \mathcal{X} , with values in the extended reals \mathbb{R}^* , which is executed on the whole set of locations L , but for the modified signals equal to $-\infty$ for locations not satisfying the metric bounds for ℓ . The function \mathcal{X} is defined below.

Definition 4.4. Given a finite set of locations L and two functions $s_1 : L \rightarrow \mathbb{R}^*$, $s_2 : L \rightarrow \mathbb{R}^*$. The function $\mathcal{X} : \mathbb{N} \times L \rightarrow \mathbb{R}$ is inductively defined as:

- (1) $\mathcal{X}(0, \ell) = s_1(\ell)$
- (2) $\mathcal{X}(i+1, \ell) = \min(\mathcal{X}(i, \ell), \min_{\ell' \in E\ell'}(\max(\mathcal{X}(i, \ell'), s_2(\ell'))))$

The algorithm then computes the function \mathcal{X} iteratively, until a fixed-point is reached.

Theorem 4.5. *Let s_1 and s_2 be as in Definition 4.4, and*

$$s(\ell) = \max_{A \subseteq L, \ell \in A} (\min(\min_{\ell' \in A} s_1(\ell'), \min_{\ell' \in B^+(A)} s_2(\ell'))),$$

then

$$\lim_{i \rightarrow \infty} \mathcal{X}(i, \ell) = s(\ell), \quad \forall \ell \in L.$$

Moreover, $\exists K > 0$ s. t. $\mathcal{X}(j, \ell) = s(\ell), \forall j \geq K$.

The following corollary provides the correctness of the method. It shows that, when \mathcal{X} is computed for the modified signals constructed by the algorithm, it returns effectively the quantitative satisfaction score of the spatial surrounded operator.

Corollary 4.6. *Given an $\hat{\ell} \in L$, let $\psi = \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2$ and*

$$s_1(\ell) = \begin{cases} \rho(\varphi_1, \mathbf{x}, t, \ell) & \text{if } 0 \leq d(\hat{\ell}, \ell) \leq d_2 \\ -\infty & \text{otherwise.} \end{cases}$$

$$s_2(\ell) = \begin{cases} \rho(\varphi_2, \mathbf{x}, t, \ell) & \text{if } d_1 \leq d(\hat{\ell}, \ell) \leq d_2 \\ -\infty & \text{otherwise.} \end{cases}$$

Then $\rho(\psi, \mathbf{x}, t, \hat{\ell}) = s(\hat{\ell}) = \max_{A \subseteq L, \hat{\ell} \in A} (\min(\min_{\ell \in A} s_1(\ell), \min_{\ell \in B^+(A)} s_2(\ell)))$.

Algorithm 2 Quantitative monitoring for the surrounded operator

```

1: inputs:  $\ell, \psi = \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2, h, T$ 
2: for all  $\ell' \in L$  do
3:   if  $0 \leq d(\ell, \ell') \leq d_2$  then
4:     compute  $s_{\varphi_1, \ell'}$ 
5:     if  $d(\ell, \ell') \geq d_1$  then
6:       compute  $s_{\varphi_2, \ell'}$ 
7:     else  $s_{\varphi_2, \ell'} = -\infty$ 
8:     else  $s_{\varphi_1, \ell'} = -\infty, s_{\varphi_2, \ell'} = -\infty$ 
9:   end for
10: for all  $t \in \{0, h, 2h, \dots, T\}$  do
11:   for all  $\ell' \in L$  do
12:      $\mathcal{X}_{prec}(\ell') = +\infty$ 
13:      $\mathcal{X}(\ell') = s_{\varphi_1, \ell'}(t)$ 
14:   end for
15:   while  $\exists \ell' \in L$ , s.t.  $\mathcal{X}_{prec}(\ell') \neq \mathcal{X}(\ell')$  do
16:      $\mathcal{X}_{prec} = \mathcal{X}$ 
17:     for all  $\ell' \in L$  do
18:        $\mathcal{X}(\ell') = \min(\mathcal{X}_{prec}(\ell'), \min_{\ell'' \in E\ell'}(\max(s_{\varphi_2, \ell''}(t), \mathcal{X}_{prec}(\ell''))))$ 
19:     end for
20:   end while
21:    $s_{\psi, \ell}(t) = \mathcal{X}(\ell)$ 
22: end for
23: return  $s_{\psi, \ell}$ 

```

In order to discuss the complexity of the monitoring procedure, we need an upper bound on the number of iterations of the algorithm. This is given by the following.

Proposition 4.7. *Let d_G be the diameter of the graph G and $\mathcal{X}(\ell)$ the fixed point of $\mathcal{X}(i, \ell)$, then $\mathcal{X}(\ell) = \mathcal{X}(d_G + 1, \ell)$ for all $\ell \in L$.*

It follows that the computational cost for each *single* location is $O(d_G|L|m)$, where m is the number of sampled time-points. The cost for *all* locations is therefore $O(d_G|L|^2m)$.

The proofs of Theorem 4.5, Corollary 4.6 and Proposition 4.7 are reported in Appendix A.2.

4.3. SSTL Monitoring Implementation. To support qualitative and quantitative monitoring of SSTL properties, a prototype tool has been developed. This tool, developed in **Java**, consists of a **Java** library (jSSTL API) and a front-end, integrated in ECLIPSE. Both the library and the ECLIPSE plugin can be downloaded from <https://github.com/Quanticol/jsstl>. All the scenarios considered in this paper are available at <https://github.com/Quanticol/jsstl-examples> to permit the replication of experiments.

The library can be used to integrate jSSTL within other applications and tools, whereas the ECLIPSE plugin provides a user friendly interface to the tool. Furthermore, the modular approach of the implementation allows us to develop different front-ends for jSSTL. The tool has been described in more detail in [37].

The ECLIPSE plugin provides a simple user interface to specify and verify SSTL properties of spatio-temporal trajectories generated from the simulation of a system or from real observations. We can specify the properties, describe a model of the space (i.e., its graph structure), import the trajectories and then verify whether such trajectories satisfy the specified properties.

In Figure 7, a snapshot is shown of the ECLIPSE plugin. It provides an *editor* for jSSTL, containing the script with the SSTL properties that we want to analyse in our scenario (on the left) and a *view* to visualise the space model, the data and the results of the analyses (on the right).

Implementation performance. Model (3.1) has been coded in Matlab/Octave, and the monitoring has been performed by our Java implementation. As time performance, the verification of property φ_{spot} took 1.04s (Boolean) and 69.39s (quantitative) for all locations and 100 time points, while property φ_{pattern} took 1.81s and 70.06s, and property φ_{pert} took 28, 19s and 55, 31s, respectively. The computation of the distance matrix can be done just once because it remains always the same for a given system, this takes about 23s. All the experiments were run on an Intel Core i5 2.6 GHz CPU, with 8GB 1600 MHz RAM.

5. SSTL ANALYSIS OF STOCHASTIC SPATIO-TEMPORAL SYSTEMS

SSTL can also be applied to describe properties of stochastic spatio-temporal systems. Stochastic models describe the evolution of systems in space and time that show noisy behaviour, due to internal random mechanisms (like in epidemic spreading models), or environmental effects (typically captured with Stochastic Differential Equations, SDEs). Independently from the mathematical device to describe them, stochastic models induce a probability measure on the space of all possible traces (i.e. on the so-called Skorokhod

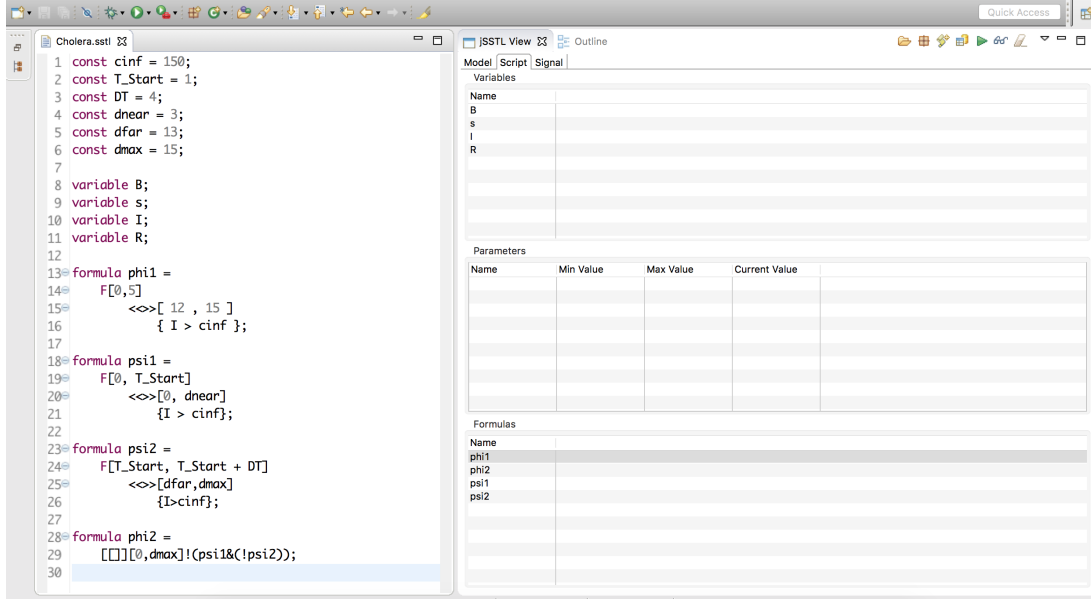


Figure 7: The jSSTL ECLIPSE plugin

space, the space of càdlàg functions, which are piecewise continuous functions from time taking real values). Each SSTL formula φ describes a subset of these traces, those satisfying it, which is measurable (with respect to the topology of the Skorokhod space), as can be proved by a simple adaptation of the proof in [6], owing to the discrete nature of space considered here. Measurability implies that we can calculate the probability $p(\varphi)$ of this set, which is known as the satisfaction probability of φ for the stochastic model considered. The goal for analysing these systems is thus to compute such satisfaction probability. Due to computational unfeasibility of exact methods [5], the mainstream approach in verification is to rely on Statistical Model Checking [40], which combines simulation of the stochastic model (i.e. an algorithm that samples traces according to the probability distribution of the model in the Skorokhod space) with a monitoring routine for the property φ . More specifically, every time a trace is generated by the simulator, it is passed to the Boolean monitor, which returns either 0 (false) or 1 (true). Probabilistically, this can be seen as a sample of a Bernoulli random variable, having probability $p(\varphi)$ of observing 1. From a finite sample of such values, we can rely on standard statistical tools to estimate $p(\varphi)$ and to compute the confidence level of such an estimate. A scheme of Statistical Model Checking (SMC) is shown in Figure 8. Examples of spatio-temporal model checking to compute the approximated probabilistic satisfaction can be found in [20, 30]

The integration of SSTL monitoring with statistical model checking opens the possibility of checking spatio-temporal properties of stochastic models, providing a powerful tool to explore complex stochastic behaviours in space and time. Moreover, a similar mechanism can be put in place for the quantitative semantics. In this case, SMC enables one to estimate the average robustness of a formula (and other moments, like the variance). The combination with SMC and the quantitative semantics has been explored earlier for STL in [6] and applied to tasks like system design and parameter synthesis [6, 9, 10]. The spatio-temporal logics SpaTeL and STLCS have also been used in conjunction with statistical spatio-temporal

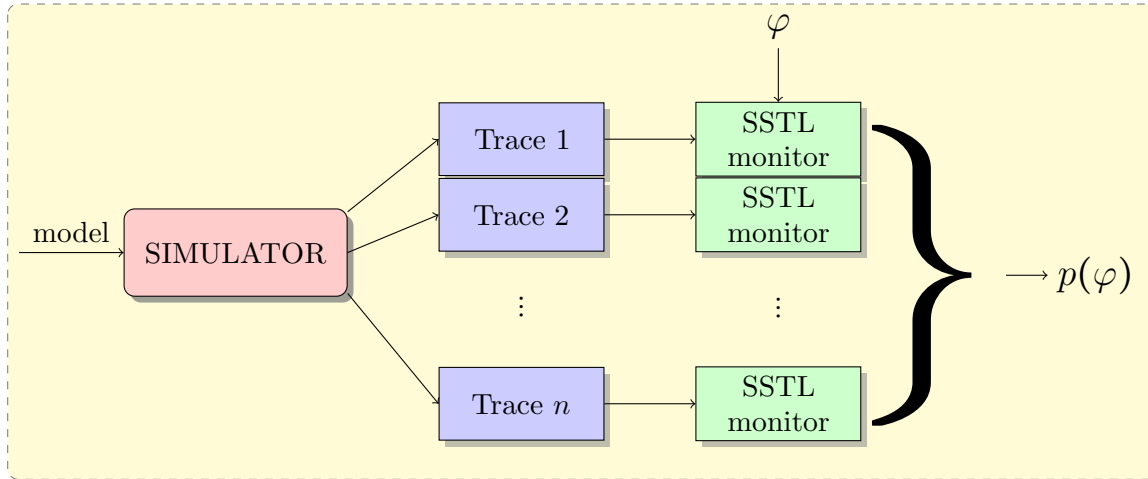


Figure 8: Scheme of Statistical Model Checking (SMC).

model checking for the Boolean semantics; see [30] and in [20], respectively, the latter focusing on usability issues in bike-sharing systems.

In the following, we give a taste of the use of SSTL to analyse stochastic models, considering a variant of the running example, in which the pattern formation mechanism is subject to external random perturbations, which is a more realistic scenario than the one described by a deterministic ODE model. An additional more complex example can be found in the next section.

Example 5.1 (Persistence of External Perturbations). We consider the effects of external perturbations of the Turing pattern formation system, adding a white Gaussian noise to the set of equations (3.1). In particular, we add a random fluctuation $\eta(\mathbf{z}, t)$, with zero-mean and covariance matrix $\langle \eta(\mathbf{z}, t), \eta(\mathbf{z}', t) \rangle = \epsilon^2 \delta z^2 \delta t$, where ϵ is the noise intensity and $\mathbf{z} = (i, j) \in [0, 32]^2$ corresponds to the position. The methodology follows [31, 32], and results in a set of Stochastic Differential Equations. In the discretisation, we set $\Delta z = 1$ and $\Delta t = 0.01$. We study how the noise affects the dynamics for fixed deterministic parameters, analysing in detail the capability of the system to maintain a specific pattern with respect to the external perturbation. To this aim, we fix the parameters of the system: $K = 32, R_1 = 1, R_2 = -12, R_3 = -1, R_4 = 16, D_1 = 5.6$ and $D_2 = 25.5$ and we evaluate formula 3.5 (the pattern formula) with parameters $h = 0.5, T_{\text{pattern}} = 19, \delta = 1, T_{\text{end}} = 30, w_1 = 1, w_2 = 6$ for different values of the noise intensity ϵ .

In Figure 9(a), we show how the satisfaction probability decreases as a function of the noise intensity ϵ ; ϵ was varied between 0 and 0.9 in steps of 0.1 units. We estimate the probability statistically from 10,000 runs for each parameter value. This result shows that an increasing intensity of noise prevents the system from establishing a regular pattern of spots of the form we have seen in a non-perturbed variant of the system. In Figure 9(b), we plot the satisfaction probability versus the average robustness degree, estimating them statistically from 10,000 runs for each parameter value. The satisfaction probability varies from 1 to 0 while the average robustness score varies from -80.01 ± 10.7 to 0.008 ± 0 . As we can see, these two quantities seem to be correlated. In other words, the higher is the




Figure 9: (a) Satisfaction probability vs. noise intensity; (b) Satisfaction probability vs. average robustness degree.

number of runs in which we find that the pattern formation property holds, the higher is the average robustness with which this happens. By varying the threshold, the Pearson's correlation coefficient between satisfaction probability and robustness degree is 0.784.

6. CASE STUDY: BIKE SHARING SYSTEM

In this section, we present an analysis of the London Santander Cycles Hire scheme, a bike sharing system, modelled as a *Population Continuous Time Markov Chain* (PCTMC) [7] with time-dependent rates. We use SSTL to study a number of spatio-temporal properties of the system and to explore their robustness considering a set of parameter values for the formulas.

6.1. Model. The Bike-Sharing System (BSS) is composed of a number of bike stations, distributed over a geographic area. Each station has a fixed number of bike slots. The users can pick up a bike, use it for a while, and then return it to another station in the area. Following [26], we model the BSS as a Population Continuous Time Markov Chain (PCTMC) with time-dependent rates, leading to a hybrid system. The model, given the bike availability in a station at time t , predicts the probability distribution of the number of available bikes in that station at time $t + h$ with $h \in [0, 40]$ minutes. The parameters of the model have been set using the historic journey data and bike availability data from January 2015 to March 2015 from the London Santander Cycles Hire scheme. In detail, we can describe the model through the following transitions:

$$\begin{aligned}
 B_i &\rightarrow S_i && \text{at rate } out_i(t) \cdot p_{out}^i, && \forall i \in (1, N) \\
 S_i &\rightarrow B_i && \text{at rate } in_i(t) \cdot \lambda_{in}^i, && \forall j \in (1, N) \\
 B_i &\rightarrow S_i + T_j^i && \text{at rate } out_i(t) \cdot p_j^i(t), && \forall i, j \in (1, N) \\
 S_j + T_j^i &\rightarrow B_j && \text{at rate } in_j^i(\#T_j^i), && \forall i, j \in (1, N).
 \end{aligned}$$

Here, B_i (respectively S_i) represents the bike agent (respectively the slot agent) in the i^{th} station, T_j^i is the bike agent travelling from pick-up station i to return station j . To these different agents we associate counts, e.g. $\#T_j^i$ denotes the population size of an agent type T_j^i , while N is the total number of stations. Each transition in the list describes a possible event changing the state of the system. In general, a transition rule like $B_i \rightarrow S_i + T_j^i$ models that an agent B_i is removed from the system (a bike leaves station i), while new agents S_i and T_j^i are added to the system (a free slot is added to station i , while the bike is set to travel towards station j). This is reflected also in the population counts: upon the firing of such a transition, $\#B_i$ will be decreased by 1, while $\#S_i$ and $\#T_j^i$ are increased by 1.

In the list of transitions above, the first transition represents the pick-up of a bike from a station i and its “removal” from the system. A bike is considered removed from the system in two cases: if its destination is a station outside the model (in case we model a subset of stations) or if the bike is taken for a time which is much longer than the usual travel time between two stations, which is about 15 minutes. This can happen, for example, if the bike is taken for a day trip. Similarly, the second transition models the return to station i of a previously removed bike. The last two transitions model the pick-up and return of a bike by users of the system that are travelling between the bike stations modelled in the system. Each transition has a rate, encoding the average frequency with which it happens, based on historic journey data. These rates, together with the updates in population counts, define a stochastic model in terms of a Continuous Time Markov Chain, which can be simulated with standard stochastic simulation algorithms. We refer the interested reader e.g. to [7] for further details. More specifically, $out_i(t)$ (respectively $in_i(t)$) is the bike pick-up (respectively return) rate in station i at time t , $\lambda_{out}^i(t)$ (respectively $\lambda_{in}^i(t)$) is the probability that a bike leaves (respectively enters from) outside the system, $p_j^i(t)$ is the probability that a journey will end at station j given that it started from station i at time t , in_j^i , instead, is the return rate of a bike pick-up in station i that will be returned in station j . The model considered is very large comprising 733 bike stations (each with 20-40 slots) and a total population of 57,713 agents (users) picking up and returning bikes. In our model the first two transitions have very small rates. They are more significant in the analysis of submodels where just a subsystem of BSS stations are considered, i.e. where there are considered stations outside the system.

6.2. Spatio-temporal analysis of the Bike Sharing Systems using SSSL. We simulate the model using Simhya [11], a Java tool for the simulation of stochastic and hybrid systems. In particular we exploit its Gibson-Bruck (GB) algorithm. Following [26], the time-dependent rates change 2 times, all at the same time unit: the first interval is 14 minutes, the second interval is 20 minutes, the third interval is 6 minutes; we simulate the model for 40 minutes. Then, we consider the trajectories only of the bike (B) and slot (S) agents, in each station. Our spatio-temporal trace is then $X_B(t, \ell) = (X_B(t, \ell), X_S(t, \ell))$, i.e. the number of bikes and free slots at each time, in each station. The space is represented by a weighted graph, where the nodes are the stations and the edges describe the connection between each station. Two nodes are connected if they are at a distance less or equal to 1 Kilometer. The weighted function $w : E \rightarrow \mathbb{R}$ returns the distance in kilometres between any two stations, where $E = L \times L$ is the set of edges and L is the set of stations.

One of the main problems of these systems consists in the availability of bikes or free slots in each station. Two important questions related to this issue from a user's point of view are:

- if I don't find a bike (free slot, resp.) is there another station at the distance less than a certain value where I can find a bike (free slot, resp.) slot?
- if I don't find a bike (free slot, resp.), how long should I wait before another user returns a bike or picks one up, respectively?

These concerns can be expressed by the SSTL properties described below. The first one is given by:

$$\varphi_1 = \mathcal{G}_{[0, T_{end}]} \{ \diamond_{[0, d]} (B > 0) \wedge \diamond_{[0, d]} (S > 0) \} \quad (6.1)$$

A station ℓ satisfies φ_1 if and only if it is always true that, between 0 and T_{end} minutes, there exists a station at a distance less than or equal to d , that has at least one bike and a station at a distance less or equal to d that has at least one free slot.

In the analysis, we explore the value of parameter $d \in [0, 1]$ kilometres to see how the satisfaction of the property changes in each location. Figure 10 shows the approximate satisfaction probability p_{φ_1} for 1000 runs for all the stations, for (a) $d = 0$, (b) $d = 0.2$, (c) $d = 0.3$ and (d) $d = 0.5$ km. We can see that for $d = 0$ many stations have a high probability to be full or empty (indicated by red points), with standard deviation of all the locations in the range $[0, 0.0158]$ and mean standard deviation 0.0053. However, increasing d to $d = 0.2$ km, i.e. allowing a search area of up to 200 metres from the station that currently has no bikes, or no slots resp., we greatly increase the satisfaction probability of φ_1 , with a standard deviation that remains in the same range and mean standard deviation of 0.0039. For $d = 0.5$, the probability of p_{φ_1} is greater than 0.5 for all the stations; standard deviation is in the range $[0, 0.0142]$ and mean stdv is 0.0002. Figure 12 (a) shows the satisfaction probability of some BBS stations vs distance $d=[0,1.0]$.

The second property “if I wait t minutes, I will always find a bike and a free slot” can instead be formalised by the following property:

$$\varphi_2 = \mathcal{G}_{[0, T_{end}]} \{ \mathcal{F}_{[0, t]} ((B > 0) \wedge (S > 0)) \} \quad (6.2)$$

A station ℓ satisfies φ_2 if and only if it is always true that, for each time $h \in [0, T_{end}]$, eventually, in a time between h and $h + t$, there will be a bike and a free slot available in ℓ .

In the analysis, we explore parameter $t \in [0, 10]$ minutes to see how the satisfaction of the property changes in each location with respect to the waiting time t for a bike or a free slot. Figure 11 shows the approximate satisfaction probability p_{φ_2} for 1000 runs for all the stations, for (a) $t = 0$ and (b) $t = 10$. In this case, we can see that waiting a certain amount of time does not greatly increase the probability to always find a bike or a free slot in the stations; even after waiting for 10 minutes Fig. 11 (b) shows that there are still many stations that satisfy φ_2 with a nearly zero probability. This means that there are stations that remain full or empty for at least 10 minutes. Considering that the average human walking speed is about 5.0 kilometres per hour (km/h), we can conclude that if users do not find a bike or a free slot in a station, they usually have more probability to find the bike/free slot in another station rather than when waiting in the same station. The standard deviation remains in the interval $[0, 0.0159]$. Figure 12 (b) shows the satisfaction probability of some BBS stations vs the time $t=[0,10]$.

Using the somewhere operator, one can only check if there is a bike in a station within a certain distance. It may be the case that the station with a bike is located in the opposite




Figure 10: Approximate satisfaction probability of formula φ_1 for 1000 runs for each BSS station for (a) $d = 0$, (b) $d = 0.2$, (c) $d = 0.3$ and (d) $d = 0.5$. The value of the probability is given by the color legend.

direction the user is headed to, costing her a large deviation and more time to pick up the bike. A more precise check is given by the property below:

$$\varphi_3 = \mathcal{G}_{[0, T_{end}]} \{ B = 0 \Rightarrow (B = 0) \mathcal{S}_{[0, d]}(B > 0) \} \quad (6.3)$$

A station ℓ satisfies φ_3 if and only if it is always true, for each time $h \in [0, T_{end}]$, that a station with no bikes is surrounded by stations with at least one bike at a distance less than or equal to d . A similar property can be defined for free slots. In other words, this would guarantee that a user who finds a station without bikes will find a station with at least one bike within distance d , no matter in which direction the user looks for another station. In the analysis, we explore parameter $d \in [0, 0.2]$ kilometers to see how the satisfaction of the property changes in each location. The results show that increasing the distance does not significantly increase the probability of φ_3 ; that is, the probability to be surrounded by stations with at least one bike does not increase. Figure 12 (c) shows the satisfaction probability of some BBS stations vs the distance $d \in [0, 0.2]$. The standard deviation remains in the interval $[0, 0.0158]$. In other words, there are always some stations with almost a constant

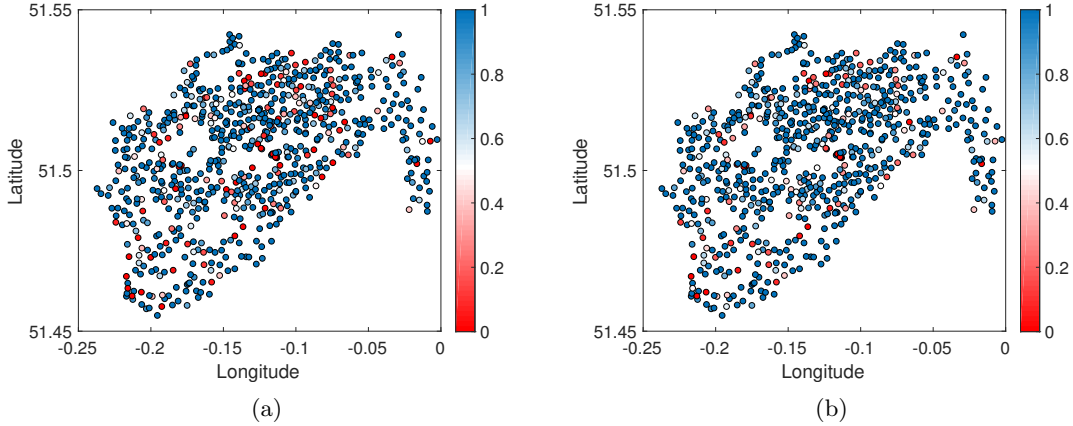


Figure 11: Approximate satisfaction probability degree of formula φ_2 for 1000 runs for each BSS station for (a) $t = 0$ and (b) $t = 10$. The value of the degree is given by the color legend.

number of bikes/free slots in some directions. This analysis suggests that considering spatial operators that include a notion of direction might be interesting for future work.

Finally, in all the previous properties we did not consider that a user will need some time to reach a nearby station. Both properties φ_1 and φ_2 can be refined to take this aspect into consideration by considering a nested spatio-temporal property:

$$\psi_1 = \mathcal{G}_{[0, T_{end}]} \{ \diamond_{[0, d]} (\mathcal{F}_{[t_w, t_w]} B > 0) \wedge \diamond_{[0, d]} (\mathcal{F}_{[t_w, t_w]} S > 0) \} \quad (6.4)$$

A station ℓ satisfies ψ_1 if and only if it is always true between 0 and T_{end} minutes that there exists a station at a distance less than or equal to d , that, eventually in a time equal to t_w (the walking time), has at least one bike and a station at a distance less than or equal to d , that, eventually in a time equal to t_w has at least one free slot.

Similarly,

$$\psi_3 = \mathcal{G}_{[0, T_{end}]} \{ B = 0 \Rightarrow (B = 0) \mathcal{S}_{[0, d]} (\mathcal{F}_{[t_w, t_w]} B > 0) \} \quad (6.5)$$

which expresses that a station ℓ satisfies ψ_3 if and only if it is always true between 0 and T_{end} minutes that if it is empty, it is surrounded by stations within a distance of at most d that have at least one bike, eventually in a walking time equal to t_w .

We consider an average walking speed of 6.0 km/h, this means for example that if we evaluate a distance $d = 0.5$ kilometers, we consider a walking time $t_w = 6$ minutes. The results of ψ_1 are very similar to the results of φ_1 . This means that there is not much difference between looking at $t = 0$ or after the walking time. This is in accordance with the result of property φ_2 where increasing the time does not significantly increase the satisfaction probability of φ_2 . This result becomes even more evident evaluating property ψ_3 . Figure 12 (d) shows the satisfaction probability of ψ_3 of some BSS stations vs distance $d = [0, 2.0]$. In this case, the probability increases slightly with respect to property φ_3 , obtaining a similar result for property φ_2 and meaning that in the selected stations (l_0 , l_{50} , l_{100} and l_{300}) the probability that a user finds a bike while walking to another station is increasing with the distance. Figure 13 shows the difference between the satisfaction probability of properties ψ_1 , φ_1 and properties ψ_3 , φ_3 for the same locations as in Figure 12. We can see that, in the

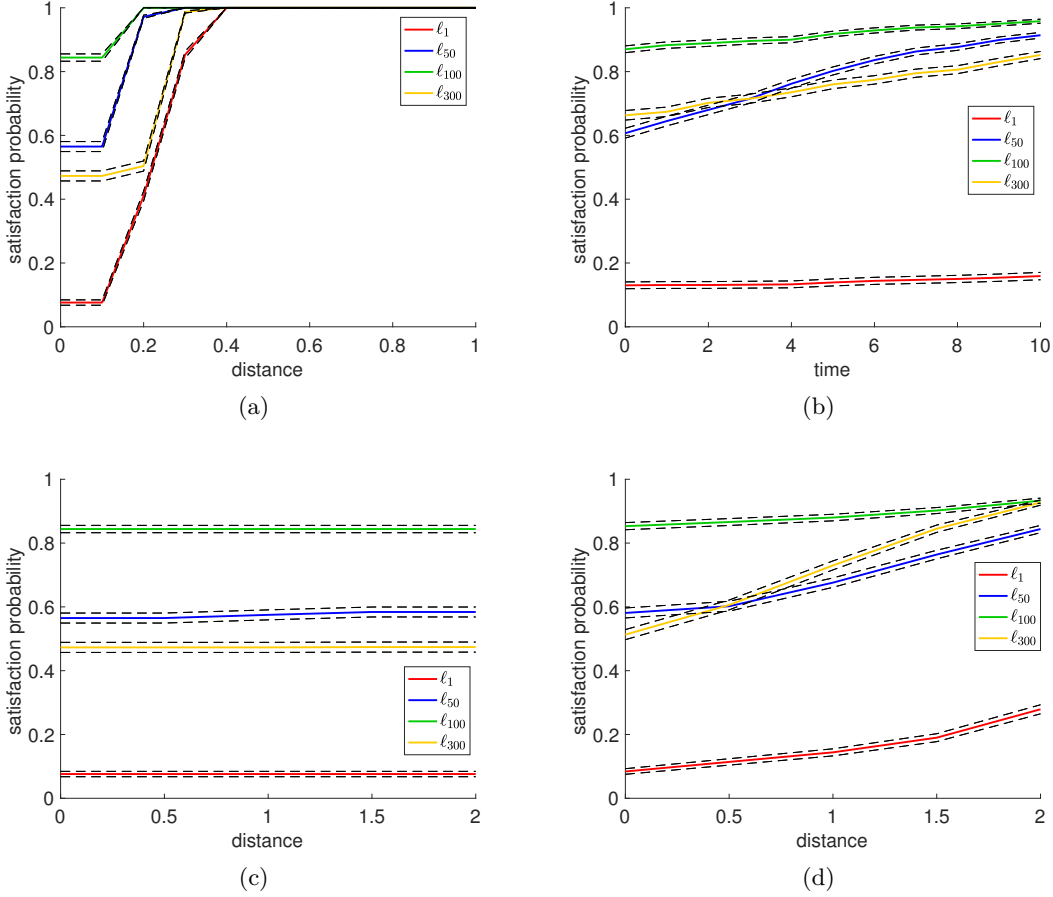


Figure 12: Approximate satisfaction probability for 1000 runs of BSS stations $\ell_1, \ell_{50}, \ell_{100}$, and ℓ_{300} for property (a) φ_1 and $d=[0,1.0]$, (b) φ_2 and $t=[0,10]$, (c) φ_3 and $d=[0,2.0]$ (d), ψ_3 and $d=[0,2.0]$.

second case, Figure 13 (b), there is a clear improvement of the satisfaction probability vs the distance parameter.

The analysis is performed using the Java implementation. As time performance, the monitoring of a single traces takes 0.207s for property φ_1 , 0.011s for property φ_2 , 0.971s for property φ_3 , 0.178s for property ψ_1 , and 0.667s for property ψ_3 . The computation of the distance matrix takes about 5.574s. All the experiments were run on an Intel Core i7 3.5 GHz CPU, with 16GB 2133 MHz RAM.

7. DISCUSSION

We defined the Signal Spatio-Temporal Logic, a spatio-temporal extension of STL [23], with two spatial operators: the somewhere operator and the spatial (bounded) surrounded operator. In SSTL spatial and temporal operators can be arbitrarily nested. We provided the logic with a Boolean and a quantitative semantics in the style of STL [23], and defined novel monitoring algorithms to evaluate such semantics on spatio-temporal trajectories. The monitoring

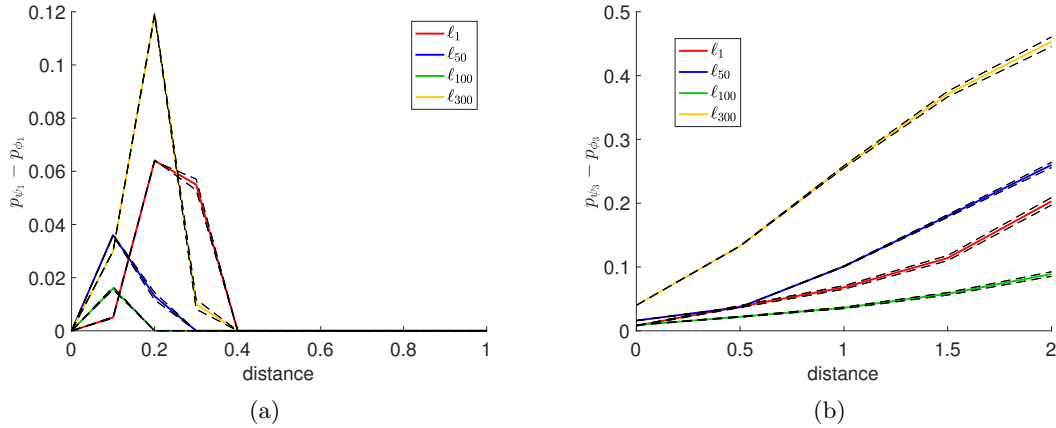


Figure 13: (a) $p_{\psi_1} - p_{\phi_1}$ vs the distance $d = [0, 1.0]$ (b) $p_{\psi_3} - p_{\phi_3}$ vs the distance $d = [0, 2.0]$.

procedures, implemented in Java, have been applied on two case studies. In the first one, we study a Turing reaction-diffusion system, modelling a process of morphogenesis [39] in which spots are formed over time. The second one is a bike sharing system modelled as a CTMC where we analyse the likelihood for users of the system to find bikes and free parking slots when they need them.

The work in this paper can be extended in several directions. First, we plan to perform a more thorough investigation of the expressivity of the logic, and to apply it on further case studies. Secondly, we plan to extend our logic to more general quasi-discrete metric spatial structures, exploiting the topological notion of closure spaces [14, 18] and extending it to the metric case. Note that the current monitoring algorithms work already for more general spatial structures, like finite directed weighted graphs, but we plan to provide a more precise characterisation of the class of discrete spatial structures on which they can be applied. We also plan to further optimise the implementation to improve performance, and additionally investigate if and how directionality can be expressed in SSTL. Finally, we plan to exploit the quantitative semantics for the robust design of spatio-temporal systems, along the lines of [6].

REFERENCES

- [1] Aiello, M., Pratt-Hartmann, I., van Benthem, J. (eds.): Handbook of Spatial Logics. Springer (2007)
- [2] Akazaki, T., Hasuo, I.: Time robustness in MTL and expressivity in hybrid system falsification. In: Proc. of CAV 2015. LNCS, vol. 9207, pp. 356–374. Springer (2015)
- [3] Alur, R., Feder, T., Henzinger, T.: The benefits of relaxing punctuality. J. ACM (1996)
- [4] Bae, K., Meseguer, J.: A rewriting-based model checker for the linear temporal logic of rewriting. Electron. Notes Theor. Comput. Sci. 290, 19–36 (Dec 2012), <http://dx.doi.org/10.1016/j.entcs.2012.11.009>
- [5] Barbot, B., Chen, T., Han, T., Katoen, J.P., Mereacre, A.: Efficient CTMC model checking of linear real-time objectives. In: Tools and Algorithms for the Construction and Analysis of Systems, pp. 128–142. Springer (2011)
- [6] Bartocci, E., Bortolussi, L., Nenzi, L., Sanguinetti, G.: System design of stochastic models using robustness of temporal properties. Theoretical Computer Science (2015)
- [7] Bortolussi, L., Hillston, J., Latella, D., Massink, M.: Continuous approximation of collective systems behaviour: a tutorial. Performance Evaluation 70(5), 317–349 (May 2013)

- [8] Bortolussi, L., Nenzi, L.: Specifying and monitoring properties of stochastic spatio-temporal systems in signal temporal logic. In: Proc. of VALUETOOLS (2014)
- [9] Bortolussi, L., Policriti, A., Silveti, S.: Logic-based multi-objective design of chemical reaction networks. In: Proc. of Hybrid Systems Biology. pp. 164–178 (2016)
- [10] Bortolussi, L., Silveti, S.: Bayesian statistical parameter synthesis for linear temporal properties of stochastic models. In: Proc. of TACAS (2018)
- [11] Bortolussi, L., Galpin, V., Hillston, J.: Hybrid performance modelling of opportunistic networks. In: Proceedings 10th Workshop on Quantitative Aspects of Programming Languages and Systems, QAPL 2012, Tallinn, Estonia, 31 March and 1 April 2012. pp. 106–121 (2012), <http://dx.doi.org/10.4204/EPTCS.85.8>
- [12] Brochenin, R., Demri, S., Lozes, E.: On the almighty wand. *Inf. Comput.* 211, 106–137 (2012), <http://dblp.uni-trier.de/db/journals/iandc/iandc211.html#BrocheninDL12>
- [13] Cardelli, L., Gardner, P., Ghelli, G.: A spatial logic for querying graphs. In: Widmayer, P., Ruiz, F.T., Bueno, R.M., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) ICALP. Lecture Notes in Computer Science, vol. 2380, pp. 597–610. Springer (2002), <http://dblp.uni-trier.de/db/conf/icalp/icalp2002.html#CardelliGG02>
- [14] Ciancia, V., Latella, D., Loretì, M., Massink, M.: Specifying and verifying properties of space. In: Proc. of IFIP-TCS (2014)
- [15] Ciancia, V., Latella, D., Loretì, M., Massink, M.: Spatial logic and spatial model checking for closure spaces. In: M. Bernardo et al. (ed.) 16th International School on Formal Methods for the Design of Computer, Communication and Software Systems: Quantitative Evaluation of Collective Adaptive System, LNCS, vol. 9700, pp. 156–201. Springer (2016)
- [16] Ciancia, V., Gilmore, S., Grilletti, G., Latella, D., Loretì, M., Massink, M.: Spatio-temporal model-checking of vehicular movement in public transport systems. *International Journal on Software Tools for Technology Transfer STTT*, Online first publication. (2018), <https://doi.org/10.1007/s10009-018-0483-8>
- [17] Ciancia, V., Grilletti, G., Latella, D., Loretì, M., Massink, M.: An experimental spatio-temporal model checker. In: Software Engineering and Formal Methods - SEFM 2015 Collocated Workshops: ATSE, HOFM, MoKMaSD, and VERY*SCART, York, UK, September 7-8, 2015, Revised Selected Papers. LNCS, vol. 9509, pp. 297–311. Springer (2015)
- [18] Ciancia, V., Latella, D., Loretì, M., Massink, M.: Model Checking Spatial Logics for Closure Spaces. *Logical Methods in Computer Science* Volume 12, Issue 4 (Oct 2016), <http://lmcs.episciences.org/2067>
- [19] Ciancia, V., Latella, D., Massink, M., Paškauskas, R.: Exploring spatio-temporal properties of bike-sharing systems. In: 2015 IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshops, SASO Workshops 2015, Cambridge, MA, USA, September 21-25, 2015. pp. 74–79. IEEE Computer Society (2015), <http://dx.doi.org/10.1109/SASOW.2015.17>
- [20] Ciancia, V., Latella, D., Massink, M., Paškauskas, R., Vandin, A.: A tool-chain for statistical spatio-temporal model checking of bike sharing systems. In: Margaria, T., Steffen, B. (eds.) Leveraging Applications of Formal Methods, Verification and Validation: Foundational Techniques - 7th International Symposium, ISoLA 2016, Imperial, Corfu, Greece, October 10-14, 2016, Proceedings, Part I. LNCS, vol. 9952, pp. 657–673. Springer (2016)
- [21] Conforti, G., Macedonio, D., Sassone, V.: Static bilog: a unifying language for spatial structures. *Fundam. Inform.* 80(1-3), 91–110 (2007), <http://dblp.uni-trier.de/db/journals/fuin/fuin80.html#ConfortiMS07>
- [22] Donzé, A., Ferrer, T., Maler, O.: Efficient robust monitoring for stl. In: Proc. of CAV (2013)
- [23] Donzé, A., Maler, O.: Robust satisfaction of temporal logic over real-valued signals. In: Proc. of FORMATS (2010)
- [24] Fainekos, G.E., Pappas, G.J.: Robustness of temporal logic specifications for continuous-time signals. *Theor. Comput. Sci.* (2009)
- [25] Fainekos, G.E., Pappas, G.J.: Robustness of temporal logic specifications for continuous-time signals. *Theoretical Computer Science* (2009)
- [26] Feng, C., Hillston, J., Reijnsbergen, D.: Moment-Based Probabilistic Prediction of Bike Availability for Bike-Sharing Systems. In: Quantitative Evaluation of Systems. pp. 139–155. Lecture Notes

- in *Computer Science*, Springer, Cham (Aug 2016), https://link.springer.com/chapter/10.1007/978-3-319-43425-4_9
- [27] Galton, A.: A generalized topological view of motion in discrete space. *Theoretical Computer Science* 305(1–3), 111 – 134 (2003), <http://www.sciencedirect.com/science/article/pii/S0304397502007016>
 - [28] Galton, A.: The mereotopology of discrete space. In: Freksa, C., Mark, D. (eds.) *Spatial Information Theory. Cognitive and Computational Foundations of Geographic Information Science. Lecture Notes in Computer Science*, Springer Berlin Heidelberg (1999)
 - [29] Gol, E.A., Bartocci, E., Belta, C.: A formal methods approach to pattern synthesis in reaction diffusion systems. In: *Proc. of CDC* (2014)
 - [30] Haghghi, I., Jones, A., Kong, J.Z., Bartocci, E., R., G., Belta, C.: SpaTeL: A Novel Spatial-Temporal Logic and Its Applications to Networked Systems. In: *Proc. of HSCC* (2015)
 - [31] Leppanen, T., Karttunen, M., Barrio, R.A., Kaski, K.: The Effect of Noise on Turing Patterns. *ResearchGate* 150, 367 (Jan 2003), https://www.researchgate.net/publication/230720532_The_Effect_of_Noise_on_Turing_Patterns
 - [32] Lesmes, F., Hochberg, D., Morán, F., Pérez-Mercader, J.: Noise-Controlled Self-Replicating Patterns. *Physical Review Letters* 91(23) (Dec 2003), <http://link.aps.org/doi/10.1103/PhysRevLett.91.238301>
 - [33] Maler, O., Nickovic, D.: Monitoring temporal properties of continuous signals. In: *Proc. FORMATS* (2004)
 - [34] Mari, L., Bertuzzo, E., Righetto, L., Casagrandi, R., Gatto, M., Rodriguez-Iturbe, I., Rinaldo, A.: Modelling cholera epidemics: the role of waterways, human mobility and sanitation. *Journal of The Royal Society Interface* (2012)
 - [35] Meseguer, J.: The temporal logic of rewriting: A gentle introduction. In: *Concurrency, Graphs and Models. Lecture Notes in Computer Science*, vol. 5065, pp. 354–382. Springer (2008), <http://dblp.uni-trier.de/db/conf/birthday/montanari2008.html#Meseguer08>
 - [36] Nenzi, L., Bortolussi, L., Ciancia, V., Loreti, M., Massink, M.: Qualitative and quantitative monitoring of spatio-temporal properties. In: *Runtime Verification (RV). LNCS*, vol. 9333, p. 21–37. Springer (2015)
 - [37] Nenzi, L., Bortolussi, L., Loreti, M.: jSSTL - a tool to monitor spatio-temporal properties. In: *VALUE-TOOLS* (2016)
 - [38] Reif, J.H., Sistla, A.: A multiprocess network logic with temporal and spatial modalities. *Journal of Computer and System Sciences* 30(1), 41–53 (February 1985)
 - [39] Turing, A.M.: The Chemical Basis of Morphogenesis. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* (1952)
 - [40] Younes, H.L.S., Kwiatkowska, M.Z., Norman, G., Parker, D.: Numerical vs. statistical probabilistic model checking: An empirical study. In: *Proc. of 2004, the 10th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, Barcelona, Spain, March 29 - April 2* (2004)

APPENDIX A. PROOFS

In this appendix, we present the proofs of Proposition 4.3 and the correctness of the Boolean (Algorithm 1) and the quantitative (Algorithm 2) monitoring algorithms for the surrounded operator.

Proposition 4.3. *Let the primary signal \mathbf{x} be Lipschitz continuous, as well as the functions defining the atomic predicates. Let M be a Lipschitz constant for all secondary signals, and h be the discretisation step. Given a SSTL formula φ , let $u(\varphi)$ count the number of temporal until operators in φ , and denote by $\rho(\varphi, \mathbf{x})$ its satisfaction score over the trace \mathbf{x} and by $\rho(\varphi, \hat{\mathbf{x}})$ the satisfaction score over the discretised version $\hat{\mathbf{x}}$ of \mathbf{x} with time step h . Then $\|\rho(\varphi, \mathbf{x}) - \rho(\varphi, \hat{\mathbf{x}})\| \leq u(\varphi)Mh$.*

Proof. We first observe that the monitoring algorithm for Boolean and spatial operators preserve the error of the input quantitative signals. This means that if $\|s_{\varphi_j, \ell} - \hat{s}_{\varphi_j, \ell}\| \leq \varepsilon$, then $\|s_{\psi, \ell} - \hat{s}_{\psi, \ell}\| \leq \varepsilon$, for ψ one of $\neg\varphi_1$, $\varphi_1 \wedge \varphi_2$, $\varphi_1 \mathcal{S}_{[w_1, w_2]} \varphi_2$, $\diamond_{[w_1, w_2]} \varphi_1$. Hence, temporal discretisation introduces errors only for temporal operators.

Now, let $I = [t_1, t_2]$ be such that $t_j = k_j h$, and denote the Minkowski sum by \oplus , so that $t \oplus I = [t + t_1, t + t_2]$. Denote by \hat{I} the discretised version of I , with step h , $\hat{I} = \{k_1 h, (k_1 + 1)h, \dots, k_2 h\}$. We observe two facts for the maximum, with identical statements holding for the minimum.

- Let $f(t)$ be Lipschitz with constant K . Let $g(t) = \max_{t' \in t \oplus I} f(t')$ and $\hat{g}(t) = \max_{t' \in t \oplus \hat{I}} f(t')$. Then $\|g(t) - \hat{g}(t)\| \leq Kh/2$. This holds by applying the Lipschitz property between a generic point in $t \oplus I$ and the closest point in $t \oplus \hat{I}$, and noting that the maximum distance between such points is $h/2$.
- If \tilde{f} is such that $\|\tilde{f}(t) - f(t)\| \leq \varepsilon$ uniformly in t , and we let g, \hat{g} as above, and $\tilde{g}(t) = \max_{t' \in t \oplus \hat{I}} \tilde{f}(t')$, then

$$\|g(t) - \tilde{g}(t)\| \leq \|g(t) - \hat{g}(t)\| + \|\hat{g}(t) - \tilde{g}(t)\| \leq Kh/2 + \varepsilon.$$

Hence, the second property implies that the additional error we introduce by evaluating a time bounded until is an additive term no larger than Kh , as in the definition of the quantitative semantics of the until, there are a nested minimum and a maximum over dense time intervals. Hence the total error will be bounded by Kh times the number of temporal operators. \square

A.1. Correctness of the Boolean Monitoring Algorithm for the Surrounded Operator. In this section we prove the correctness of Algorithm 1; let us call the algorithm BoolSurround.

Theorem 4.2. *Given a graph $G = (L, w, E)$, two properties φ_1 and φ_2 , a trace \mathbf{x} and a location ℓ , let $s_{\psi, \ell} = \text{BoolSurround}(G, \mathbf{x}, \varphi_1, \varphi_2, \ell)$ and $\mathcal{I}_{s_{\psi, \ell}}$ be the minimal interval covering consistent with $\{s_{\varphi_1, \ell'}, s_{\varphi_2, \ell'}\}_{\ell' \in L_{[0, w_2]}^\ell}$, then, for all $I_i \in \mathcal{I}_{s_{\psi, \ell}}$*

$$s_{\psi, \ell}(I_i) = 1 \iff (\mathbf{x}, t, \ell) \models \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2 \quad \forall t \in I_i$$

Proof. First we note that $s_{\psi, \ell}(I_i) = 1 \iff \ell \in V_{I_i}$, where V_{I_i} is the set V at the end of the iteration of the I_i interval. Then, it is enough to prove that, for all $I_i \in \mathcal{I}_{s_{\psi, \ell}}$

$$\ell \in V_{I_i} \iff (\mathbf{x}, t, \ell) \models \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2 \quad \forall t \in I_i.$$

Furthermore, for the definition of the minimal interval covering (Definition 4.1), $s_{\varphi_1, \ell'}, s_{\varphi_2, \ell'}, \ell' \in L_{[0, w_2]}^\ell$ have the same value in each $I_i \in \mathcal{I}_{s_{\psi, \ell}}$. This implies, for the Boolean semantics of the surrounded operator, that, $(\mathbf{x}, \hat{t}, \ell) \models \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2$ for a specific $\hat{t} \in I_i$ if and only if it satisfies the property for all $t \in I_i$.

Let's consider now the distance constraints of the formula. We redefine the property φ_1 and φ_2 in this way:

$$(\mathbf{x}, t, \ell) \models \hat{\varphi}_1 \iff (\mathbf{x}, t, \ell') \models \varphi_1 \wedge d(\ell, \ell') \leq d_2,$$

and

$$(\mathbf{x}, t, \ell) \models \hat{\varphi}_2 \iff (\mathbf{x}, t, \ell') \models \varphi_1 \wedge d(\ell, \ell') \in [d_1, d_2].$$

Hence, we have that $V = \{\ell' \mid s_{\hat{\varphi}_1, \ell'}(I_i) = 1\}$ and $Q = \{\ell' \mid s_{\hat{\varphi}_2, \ell'}(I_i) = 1\}$.

Furthermore, a location ℓ is a bad location if it can reach a point satisfying $\neg\hat{\varphi}_1$ passing for a node $\neg\hat{\varphi}_2$. Let's consider the set

$$\mathcal{C}_\ell = \{i \in \mathbb{N} \mid \exists p: \ell \rightsquigarrow \infty.G, (\mathbf{x}, t, p(i)) \models \neg\hat{\varphi}_1, \text{ and } \forall j \in \{1, \dots, i\} (\mathbf{x}, t, p(j)) \models \neg\hat{\varphi}_2\},$$

where $p: \ell \rightsquigarrow \infty.G$ is a path of the graph G , starting from ℓ , then

$$(\mathbf{x}, t, \ell) \models \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2 \iff (\mathbf{x}, t, \ell) \models \hat{\varphi}_1 \wedge \mathcal{C}_\ell = \emptyset.$$

Hence, what we have to prove at the end is that

$$\ell \in V_{I_i} \iff (\mathbf{x}, t, \ell) \models \hat{\varphi}_1 \wedge \mathcal{C}_\ell = \emptyset, \text{ for a } t \in I_i.$$

We will prove it by induction. From this point on, we fix the trace \mathbf{x} and the time t and we will write $\ell \models \varphi$ to indicate $(\mathbf{x}, t, \ell) \models \varphi$ and V for V_{I_i} .

(\Rightarrow): We have to prove that if $(\mathbf{x}, t, \ell) \models \hat{\varphi}_1 \wedge \min \mathcal{C}_\ell = k$ then ℓ is removed at iteration k from V .

(basis step): $\ell \models \hat{\varphi}_1 \wedge \min \mathcal{C}_\ell = 1$ then $p(1) \models \neg\hat{\varphi}_1 \wedge p(1) \models \neg\hat{\varphi}_2$. This implies that $\exists \ell' \in T = B^+(Q \cup V)$, and $(\ell, \ell') \in E$, then ℓ is removed from V at the first iteration.

(inductive step): Let's suppose that if $\ell \models \hat{\varphi}_1 \wedge \min \mathcal{C}_\ell = k$ then ℓ is removed at iteration k from V . We have to prove that this is true also for $k+1$. Let's suppose that $\ell \models \hat{\varphi}_1 \wedge \min \mathcal{C}_\ell = k+1$. This implies that $p(k+1) \models \neg\hat{\varphi}_1$ and $\forall j \in \{1, \dots, k\}, p(j) \models \neg\hat{\varphi}_2$. But if $k+1 = \min \mathcal{C}_\ell$ then $\ell' = p(1) \models \hat{\varphi}_1$ and $\min \mathcal{C}_{\ell'} = k$, i.e., ℓ' is removed at iteration k from V_{I_i} , then ℓ is removed at iteration $k+1$ because $(\ell, \ell') \in E$.

(\Leftarrow): We have to prove that if ℓ is removed at iteration k from V_{I_i} then $\ell \models \hat{\varphi}_1 \wedge \min \mathcal{C}_\ell = k$.

(basis step): If ℓ is removed from V at the first iteration then $\exists \ell' \in T$ s.t. $(\ell, \ell') \in E$ and $(\mathbf{x}, t, \ell') \models \neg\hat{\varphi}_1 \wedge \neg\hat{\varphi}_2$, this implies $\min \mathcal{C}_\ell = 1$.

(inductive step): Let's suppose that if ℓ is removed at iteration k from V then $\ell \models \hat{\varphi}_1 \wedge \min \mathcal{C}_\ell = k$. We have to prove that this is true also for $k+1$. Let's suppose that ℓ is removed at iteration $k+1$ from V . This implies that $\exists \ell' \in L$ s.t. $(\ell, \ell') \in E$ and $\ell' \in T$ but this means that ℓ' was removed from V at the previous iteration k and from the inductive step we have $\min \mathcal{C}_{\ell'} = k$. If $\min \mathcal{C}_{\ell'} = k$ then $\exists p: \ell' \rightsquigarrow \infty.G$ s.t. $p(k) \models \neg\hat{\varphi}_1$ and, $\forall i \in \{1, \dots, k\}, p(i) \models \neg\hat{\varphi}_2$. But $(\ell, \ell') \in E$ and $(\mathbf{x}, t, \ell') \models \neg\hat{\varphi}_2$ (because $\ell' \in T$) then $\exists p': \ell \rightsquigarrow \infty.G$ s.t. $p'(k+1) \models \neg\hat{\varphi}_1$ and, $\forall i \in \{1, \dots, k+1\}, p'(i) \models \neg\hat{\varphi}_2$. This implies that $\min \mathcal{C}_\ell \leq k+1$, but it can be less than $k+1$ because in that case it has to be removed before. Hence, we can conclude that $\min \mathcal{C}_\ell = k+1$. \square

A.2. Correctness of the Quantitative Monitoring Algorithm for the Surrounded Operator. In this section, we present the proofs of Theorem 4.5, Corollary 4.6 and Proposition 4.7. For simplicity, we report again the statements.

Theorem 4.5. *Let s_1 and s_2 be as in Definition 4.4, and*

$$s(\ell) = \max_{A \in L, \ell \in A} (\min_{\ell' \in A} (\min s_1(\ell'), \min_{\ell' \in B^+(A)} s_2(\ell'))),$$

then

$$\lim_{i \rightarrow \infty} \mathcal{X}(i, \ell) = s(\ell), \quad \forall \ell \in L.$$

Moreover, $\exists K > 0$ s. t. $\mathcal{X}(j, \ell) = s(\ell), \forall j \geq K$.

Note that s is equivalent to the quantitative semantics of the surrounded operator $\varphi_1 \mathcal{S} \varphi_2$, with s_i denoting the robustness of φ_i , without the distance constraints. We first present two lemmas, followed by the proof of Theorem 4.5.

Lemma A.1. *If $\mathcal{X}(k+1, \ell) = \mathcal{X}(k, \ell)$ for all $\ell \in L$ then, $\forall i > k, \mathcal{X}(i, \ell) = \mathcal{X}(k, \ell)$.*

Proof. By induction.

(basis step): $i=k+1$ is true by hypothesis,

(inductive step): suppose the assert holds for $i > k$, i.e., $\mathcal{X}(i, \ell) = \mathcal{X}(k, \ell)$ (I.H.), then we have to prove that it holds for $i+1$.

$$\begin{aligned} \mathcal{X}(i+1, \ell) &= \min(\mathcal{X}(i, \ell), \min_{\ell' | \ell E \ell'} (\max(\mathcal{X}(i, \ell'), s_2(\ell')))) && \{\text{by Def. of } \mathcal{X}\} \\ &= \min(\mathcal{X}(k, \ell), \min_{\ell' | \ell E \ell'} (\max(\mathcal{X}(k, \ell'), s_2(\ell')))) && \{\text{by I.H.}\} \\ &= \mathcal{X}(k+1, \ell) = \mathcal{X}(k, \ell). && \{\text{by Def. of } \mathcal{X}\} \quad \square \end{aligned}$$

Lemma A.2. *Let A_ℓ be the subregion that maximizes $s(\ell)$, then, $\forall \ell' \in A_\ell, s(\ell') \geq s(\ell)$.*

Proof. If A_ℓ is the subregion that maximizes $s(\ell)$ then

$$s(\ell) = \min(\min_{\ell' \in A_\ell} s_1(\ell'), \min_{\ell' \in B^+(A_\ell)} s_2(\ell'))$$

Suppose by contradiction that $\exists \hat{\ell} \in A_\ell$ s.t. $s(\hat{\ell}) < s(\ell)$. Let $Q = \{A \in L, \hat{\ell} \in A\}$. Then

$$s(\hat{\ell}) = \max_{A \in Q} (\min_{\ell' \in A} (\min s_1(\ell'), \min_{\ell' \in B^+(A)} s_2(\ell')))$$

and $s(\hat{\ell}) < s(\ell)$ implies

$$\max_{A \in Q} (\min_{\ell' \in A} (\min s_1(\ell'), \min_{\ell' \in B^+(A)} s_2(\ell'))) < \min(\min_{\ell' \in A_\ell} s_1(\ell'), \min_{\ell' \in B^+(A_\ell)} s_2(\ell'))$$

But A_ℓ is a subset of L and $\hat{\ell} \in A_\ell$ therefore $A_\ell \in Q$, thus the inequality can not hold. \square

Proof of Theorem 4.5. We have to prove that (1) $\mathcal{X}(i, \ell)$ converges in a finite number of steps, in each location ℓ , to $\mathcal{X}(\ell) \in \mathbb{R}^*$ and that (2) $\forall \ell \in L, \mathcal{X}(\ell) = s(\ell)$.

(1) Convergence of \mathcal{X} .

First note that $\mathcal{X}(i, \ell) \geq \min(\mathcal{X}(i, \ell), \min_{\ell' | \ell E \ell'} (\max(\mathcal{X}(i, \ell'), s_2(\ell')))) = \mathcal{X}(i+1, \ell)$, thus $\mathcal{X}|_\ell$ is a monotonic decreasing function. Second, note that $\mathcal{X}(i, \ell) \in \{s_j(\ell) | j \in \{1, 2\}, \ell \in L\}$ is a finite set of sortable values. So, in every step, \mathcal{X} takes a value of a sortable finite set. Finally, if it happens that for a step, for all $\ell \in L, \mathcal{X}(i, \ell)$ does not change then, from

Lemma A.1, it will remain the same for all the next steps. The convergence of \mathcal{X} to the maximum fixed point follows then from Tarsky's theorem.

(2) We have to prove that $\forall \ell, \mathcal{X}(\ell) = s(\ell)$.

Let A_ℓ be the subregion that maximizes $s(\ell)$ then

$$s(\ell) = \min(\min_{\ell' \in A_\ell} s_1(\ell'), \min_{\ell' \in B^+(A_\ell)} s_2(\ell')).$$

First we prove that (2a) $\forall \ell, \mathcal{X}(\ell) \geq s(\ell)$ and then that (2b) they are equal.

(2a) To prove that $\mathcal{X}(\ell) \geq s(\ell)$ it suffices to prove that, for a generic $\ell, \forall i \in \mathbb{N}, \mathcal{X}(i, \ell) \geq s(\ell)$, and for the convergence of \mathcal{X} that $\exists j \in \mathbb{N}$ s.t. $\mathcal{X}(\ell) = \mathcal{X}(j, \ell), \forall \ell, \forall j \geq i$. The proof is by induction.

• (basis step)

$$\begin{aligned} \mathcal{X}(0, \ell) &= s_1(\ell) && \{\text{by Def. of } \mathcal{X}\} \\ &\geq \min_{\ell' \in A_\ell} s_1(\ell') && \{\text{Because } \ell \in A_\ell\} \\ &\geq \min(\min_{\ell' \in A_\ell} s_1(\ell'), \min_{\ell' \in B^+(A_\ell)} s_2(\ell')) && \{\text{min property}\} \\ &= s(\ell) && \{\text{by Def. of } s(\ell)\} \end{aligned}$$

• (inductive step) Assume $\mathcal{X}(i, \ell) \geq s(\ell)$, to prove that $\mathcal{X}(i+1, \ell) \geq s(\ell)$;

$$\mathcal{X}(i+1, \ell) = \min(\mathcal{X}(i, \ell), \min_{\ell' | \ell E \ell'} (\max(\mathcal{X}(i, \ell'), s_2(\ell')))) \quad \{\text{by Def. of } \mathcal{X}\}$$

We know by I.H. that $\mathcal{X}(i, \ell) \geq s(\ell)$, so it remains to be shown that also:

$$\min_{\ell' | \ell E \ell'} (\max(\mathcal{X}(i, \ell'), s_2(\ell'))) \geq s(\ell) \quad (\text{A.1})$$

Note that it is assumed that $\ell \in A_\ell$ and that ℓ' are direct neighbours of ℓ . Therefore we can distinguish the following two cases:

- Suppose $\ell' \in A_\ell$. By I.H. we know that $\mathcal{X}(i, \ell') \geq s(\ell')$ and by Lemma A.2 we also know that $s(\ell') \geq s(\ell)$. For what concerns $s_2(\ell')$, if $s_2(\ell') \leq \mathcal{X}(i, \ell')$ then the max leads to $\mathcal{X}(i, \ell') \geq s(\ell)$. If instead $s_2(\ell') \geq \mathcal{X}(i, \ell') \geq s(\ell)$, then obviously also $s_2(\ell') \geq s(\ell)$. So inequation (A.1) holds in this case.
- Suppose $\ell' \in B^+(A_\ell)$. Then, by definition of $s(\ell)$ we know that $s_2(\ell') \geq s(\ell)$. So, if $s_2(\ell') \geq \mathcal{X}(i, \ell')$ then the inequation holds. If $\mathcal{X}(i, \ell') \geq s_2(\ell')$ then since $s_2(\ell') \geq s(\ell)$, inequation (A.1) also holds.

(2b) Suppose by contradiction that $\exists \hat{\ell} \in L$ s.t. $\mathcal{X}(\hat{\ell}) > s(\hat{\ell})$. At the fixed point we have that

$$\mathcal{X}(\hat{\ell}) = \min(\mathcal{X}(\hat{\ell}), \min_{\ell | \hat{\ell} E \ell} (\max(\mathcal{X}(\ell), s_2(\ell))))$$

This means that the inequality

$$\min_{\ell | \hat{\ell} E \ell} (\max(\mathcal{X}(\ell), s_2(\ell))) > s(\hat{\ell}) \quad (\text{A.2})$$

has to be true.

Let $A \subseteq L$, we define:

- $C(A) := \{\ell \in L | \exists \ell' \in A \text{ s.t. } \ell' E \ell \wedge \mathcal{X}(\ell) \geq s_2(\ell)\}$
- $C^i(A) = C(C^{i-1}(A))$

We can then compute the closure of C , as $C^*(A) = A \cup_{i=0}^{\infty} C^i(A)$. Because of the definition of C and the inequality (A.2), we have that $s_1(\ell) \geq \mathcal{X}(\ell) > s(\hat{\ell})$, $\forall \ell \in C^*(\{\hat{\ell}\})$, and that $s_2(\ell) > s(\hat{\ell})$, $\forall \ell \in B^+(C^*(\{\hat{\ell}\}))$; hence

$$\min\left(\min_{\ell \in C^*(\{\hat{\ell}\})} s_1(\ell), \min_{\ell \in B^+(C^*(\{\hat{\ell}\}))} s_2(\ell)\right) > s(\hat{\ell})$$

i.e.

$$\min\left(\min_{\ell \in C^*(\{\hat{\ell}\})} s_1(\ell), \min_{\ell \in B^+(C^*(\{\hat{\ell}\}))} s_2(\ell)\right) > \min\left(\min_{\ell \in A_{\hat{\ell}}} s_1(\ell), \min_{\ell' \in B^+(A_{\hat{\ell}})} s_2(\ell')\right)$$

but this contradicts the assumption of maximality of $A_{\hat{\ell}}$. \square

Corollary 4.6. *Given an $\hat{\ell} \in L$, let $\psi = \varphi_1 \mathcal{S}_{[d_1, d_2]} \varphi_2$ and*

$$s_1(\ell) = \begin{cases} \rho(\varphi_1, \mathbf{x}, t, \ell) & \text{if } 0 \leq d(\hat{\ell}, \ell) \leq d_2 \\ -\infty & \text{otherwise.} \end{cases}$$

$$s_2(\ell) = \begin{cases} \rho(\varphi_2, \mathbf{x}, t, \ell) & \text{if } d_1 \leq d(\hat{\ell}, \ell) \leq d_2 \\ -\infty & \text{otherwise.} \end{cases}$$

Then $\rho(\psi, \mathbf{x}, t, \hat{\ell}) = s(\hat{\ell}) = \max_{A \subseteq L, \hat{\ell} \in A} (\min(\min_{\ell \in A} s_1(\ell), \min_{\ell \in B^+(A)} s_2(\ell)))$.

Proof. We recall that

$$\rho(\psi, \mathbf{x}, t, \hat{\ell}) = \max_{A \subseteq L_{[0, d_2]}^{\hat{\ell}}, \ell \in A, B^+(A) \subseteq L_{[d_1, d_2]}^{\hat{\ell}}} (\min(\min_{\ell \in A} \rho(\varphi_1, \mathbf{x}, t, \ell), \min_{\ell' \in B^+(A)} \rho(\varphi_2, \mathbf{x}, t, \ell'))),$$

where $L_{[d_1, d_2]}^{\hat{\ell}} := \{\ell \in A \mid d_1 \leq d(\ell, \hat{\ell}) \leq d_2\}$. This means that $\ell \in A$ iff $d(\ell, \hat{\ell}) \leq d_2$ and, for all $\ell' \in A$, $d_1 \leq d(\ell', \hat{\ell}) \leq d_2$.

So, we consider a restricted number of subsets of L for ρ and all the possible subsets of L for s . Furthermore, the value of the locations considered by both are always the same, i.e., the value of s_1 and s_2 differ only in the locations considered by s and not by ρ . For this reason $s(\ell) \geq \rho(\ell)$.

Let A_{ρ} be the subset that maximizes ρ of $\hat{\ell}$ and A_s the subset that maximizes s of $\hat{\ell}$. And suppose by contradiction that

$$\min\left(\min_{\ell \in A_s} s_1(\ell), \min_{\ell' \in B^+(A_s)} s_2(\ell')\right) > \min\left(\min_{\ell \in A_{\rho}} \rho(\varphi_1, \mathbf{x}, t, \ell), \min_{\ell' \in B^+(A_{\rho})} \rho(\varphi_2, \mathbf{x}, t, \ell')\right),$$

but the values considered by s and not by ρ are all equal to $-\infty$ (see line 8 of Alg. 2), so if A_s has a location that cannot be considered by ρ it means that

$$\min\left(\min_{\ell \in A_s} s_1(\ell), \min_{\ell' \in B^+(A_s)} s_2(\ell')\right) = -\infty$$

but minus infinity cannot be bigger than any number. \square

Proposition 4.7. *Let d_G be the diameter of the graph G and $\mathcal{X}(\ell)$ the fixed point of $\mathcal{X}(i, \ell)$, then $\mathcal{X}(\ell) = \mathcal{X}(d_G + 1, \ell)$ for all $\ell \in L$.*

Proof. The graph diameter of G is equal to $d_g = \max_{\ell, \ell' \in L} d(\ell, \ell')$. Recall that $\mathcal{X}(d_g, \ell) \in \{s_j(\ell) \mid j \in \{1, 2\}, \ell \in L\}$ is a finite set of sortable values. At step zero the value of \mathcal{X} is equal to s_1 in all the locations. At each next step, the value of $\mathcal{X}(i, \ell)$ depends only on the value of \mathcal{X} in the same location at the previous step and the value of s_2 and \mathcal{X} in the previous

step in the direct neighbours of ℓ , $\ell' E \ell$. This means that, after a number of steps equal to the diameter of the graph, i.e., the longest shortest path of the network, \mathcal{X} , for all nodes ℓ , has taken into account the values s_1 and s_2 of all the nodes. \square