

# A Novel Spatial-Temporal Specification-Based Monitoring System for Smart Cities

Meiyi Ma, Ezio Bartocci, Eli Lifland, John Stankovic, *Life Fellow, IEEE*, and Lu Feng, *Member, IEEE*

**Abstract**—With the development of the Internet of Things, millions of sensors are being deployed in cities to collect real-time data. This leads to a need for checking city states against city requirements at runtime. In this paper, we develop a novel spatial-temporal specification-based monitoring system for smart cities. We first describe a study of over 1,000 smart city requirements, some of which cannot be specified using existing logic such as Signal Temporal Logic (STL) and its variants. To tackle this limitation, we develop SaSTL—a novel Spatial Aggregation Signal Temporal Logic—for the efficient runtime monitoring of safety and performance requirements in smart cities. We develop two new logical operators in SaSTL to augment STL for expressing spatial aggregation and spatial counting characteristics that are commonly found in real city requirements. We define Boolean and quantitative semantics for SaSTL in support of the analysis of city performance across different periods and locations. We also develop efficient monitoring algorithms that can check a SaSTL requirement in parallel over multiple data streams (e.g., generated by multiple sensors distributed spatially in a city). Additionally, we build a SaSTL-based monitoring tool to support decision making of different stakeholders to specify and runtime monitor their requirements in smart cities. We evaluate our SaSTL monitor by applying it to three case studies with large-scale real city sensing data (e.g., up to 10,000 sensors in one study). The results show that SaSTL has a much higher coverage expressiveness than other spatial-temporal logics, and with a significant reduction of computation time for monitoring requirements. We also demonstrate that the SaSTL monitor improves the safety and performance of smart cities via simulated experiments.

**Index Terms**—Signal Temporal Logic, Runtime Verification, Smart Cities.

## I. INTRODUCTION

Smart cities are emerging around the world. Examples include Chicago’s Array of Things project [1], IBM’s Rio de Janeiro Operations Center [2] and Cisco’s Smart+Connected Operations Center [3], just to name a few. Smart cities utilize a vast amount of data and smart services to enhance the safety, efficiency, and performance of city operations [4]. There is a need for monitoring city states in real-time to ensure safety and performance requirements [5]. If a requirement violation is detected by the monitor, the city operators and smart service providers can take actions to change the states, such as improving traffic performance, rejecting unsafe actions, sending alarms to police, etc. The key **challenges** of developing

such a monitor include how to use an expressive, machine-understandable language to specify smart city requirements, and how to efficiently monitor requirements that may involve multiple sensor data streams (e.g., some requirements are concerned with thousands of sensors in a smart city).

Previous works [6], [7], [8], [9] have proposed solutions to monitor smart cities using formal specification languages and their monitoring machinery. One of the latest works, CityResolver [10] uses Signal Temporal Logic (STL) [11] to support the specification-based monitoring of safety and performance requirements of smart cities. However, STL is not expressive enough to specify smart city requirements concerning *spatial* information such as “*the average noise level within 1 km of all elementary schools should always be less than 50 dB*”. There are some existing spatial extensions of STL (e.g., SSTL [12], SpaTeL [9] and STREL [13], [14], see [15] for a recent tutorial), which can express requirements such as “*there should be no traffic congestion on all the roads in the northeast direction*”. But they are not expressive enough to specify requirements like “*there should be no traffic congestion on all the roads on average*”, or “*on 90% of the roads*”, which require the aggregation and counting of signals in the spatial domain. To tackle these challenges and limitations, we develop a novel Spatial Aggregation Signal Temporal Logic (SaSTL), which extends STL with two new logical operators for expressing spatial aggregation and spatial counting characteristics which we demonstrate are commonly found in real city requirements. More specifically, this paper has the following major **contributions**:

(1) To the best of our knowledge, this is the first work studying and annotating over 1,000 real smart city requirements across different service domains to identify the gap of expressing smart city requirements with existing formal specification languages. As a result, we found that aggregation and counting signals in the spatial domain (e.g., for representing sensor signals distributed spatially in a smart city) are extremely important for specifying and monitoring city requirements.

(2) Drawing on the insights from our requirements study, we develop a new specification language SaSTL, which extends STL with a *spatial aggregation* operator and a *spatial counting* operator. SaSTL can be used to specify Point of Interests (PoIs), the physical distance, spatial relations of the PoIs and sensors, aggregation of signals over locations, degree/percentage of satisfaction and the temporal elements in a very flexible spatial-temporal scale. We define Boolean and quantitative semantics with theoretical proofs.

(3) We compare SaSTL with some existing specification

Corresponding Author: Meiyi Ma, E-mail: meiyi@virginia.edu.

Meiyi Ma, Eli Lifland, John Stankovic and Lu Feng are with the Department of Computer Science, University of Virginia, VA, 22904.

Ezio Bartocci is with the Faculty of Informatics, TU Wien, Austria.

Copyright (c) 2021 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

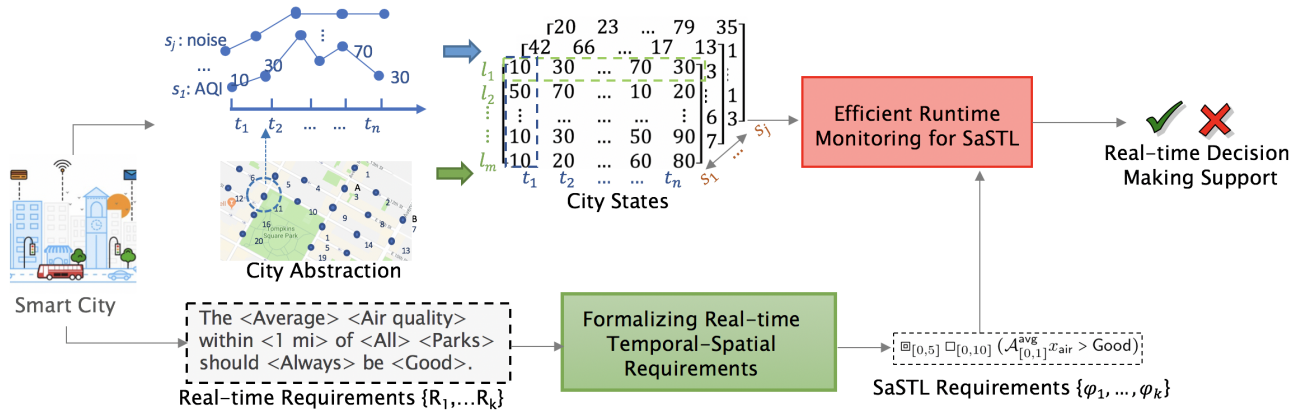


Figure 1: A framework for runtime monitoring of real-time city requirements

languages and show that SaSTL has a much higher coverage expressiveness (95%) than STL (18.4%), SSTL (43.1%) or STREL (43.1%) over 1,000 real city requirements.

(4) We develop novel and efficient monitoring algorithms for SaSTL. In particular, we present two new methods to speed up the monitoring performance: (i) dynamically prioritizing the monitoring based on cost functions assigned to nodes of the syntax tree, and (ii) parallelizing the monitoring of spatial operators among multiple locations and/or sensors.

(5) We evaluate the SaSTL monitor by applying it to monitoring real city data collected from Chicago and Aarhus. The results show that SaSTL monitor has the potential to help identify safety violations and support the city managers and citizens to make decisions. We also evaluate the SaSTL monitor on a third case study of conflict detection and resolution among smart services in simulated New York City with large-scale real sensing data (e.g., up to 10,000 sensors used in one requirement). Results of our simulated experiments show that SaSTL monitor can help improve the city's performance (e.g., 21.1% on the environment and 16.6% on public safety), with a significant reduction of computation time compared with previous approaches.

(6) We develop a SaSTL monitoring tool that can support decision making of different stakeholders in smart cities. The tool allows users (e.g., city decision maker, citizens) without any formal method background to specify city requirements and monitor city performance easily.

This paper is an extended version of [16]. We extend with the following new contributions. First, we add new quantitative semantics and monitoring algorithms, with new proofs of soundness and correctness in Section III. Compared to the conference version (Boolean semantics), the new quantitative semantics presents the monitoring results with real values, and better supports decision-makers to compare the satisfaction/violation degrees between different options. Second, we develop new monitoring algorithms for the proposed quantitative semantics and improve the monitoring algorithms for the new spatial operators in Section IV. Third, we develop a monitoring tool to support monitoring and decision making using SaSTL in smart cities in Section VI. The tool also provides a way for non-expert users to input requirements in the English language. Then the tool translates the requirements

to SaSTL formal specification automatically for monitoring. Fourth, we extend the evaluation with a new city scenario using real data from Aarhus, Denmark in Section VII. The results show that the SaSTL monitor has the potential to help identify safety violations and support city managers and citizens to make decisions. Last, we elaborate with more discussions on how to apply the SaSTL monitor in smart cities and extend the related work.

## II. APPROACH OVERVIEW

Figure 1 shows an overview of our SaSTL runtime monitoring framework for smart cities. We envision that such a framework would operate in a smart city's central control center (e.g., IBM's Rio de Janeiro Operations Center [2] or Cisco's Smart+Connected Operations Center [3]) where sensor data about city states across various locations are available in real time. The framework would monitor city states and check them against a set of smart city requirements at runtime. The monitoring results would be presented to city managers to support decision making. The framework makes abstractions of city states in the following way. The framework formalizes a set of smart city requirements (See Section III) to some machine checkable SaSTL formulas (See Section IV). Different data streams (e.g. CO emission, noise level) over temporal and spatial domains can be viewed as a 3-dimensional matrix. For any signal  $s_j$  in signal domain  $S$ , each row is a time-series data at one location and each column is a set of data streams from all locations at one time. Next, the efficient real-time monitoring for SaSTL verifies the states with the requirements and outputs the Boolean satisfaction to the decision makers, who would take actions to resolve the violation. To support decision making in real time, we improve the efficiency of the monitoring algorithm in Section V. We implement SaSTL runtime monitoring tool following this framework for city experts without any formal methods background (see Section VI). We describe more details of the framework in the following sections.

## III. ANALYSIS OF REAL CITY REQUIREMENTS

To better understand real city requirements, we conduct a requirement study. We collect and statistically analyze 1000

Table I: Examples of city requirements from different domains

Domain	Example
Transportation	Limits vehicle idling to one minute adjacent to any school, pre-K to 12th grade, public or private, in the City of New York [17].
	The engine, power and exhaust mechanism of each motor vehicle shall be equipped, adjusted and operated to prevent the escape of a trail of visible fumes or smoke for more than ten (10) consecutive seconds [18].
	Prohibit sight-seeing buses from using all bus lanes between the hours of 7:00 a.m. and 10:00 a.m. on weekdays [19].
Energy	Operate the system to maintain zone temperatures down to 55°F or up to 85°F [20].
	The total leakage shall be less than or equal to 4 cubic feet per minute per 100 square feet of conditioned floor area [21].
Environment	LA Sec 111.03 minimum ambient noise level table: ZONE M2 and M3 – DAY : 65 dB(A) NIGHT : 65 dB(A) [22].
	The total amount of HCHO emission should be less than 0.1mg per m <sup>3</sup> within an hour, and the total amount of PM10 emission should be less than 0.15 mg per m <sup>3</sup> within 24 hours [23].
Emergency	NYC Authorized emergency vehicles may disregard 4 primary rules regarding traffic [24].
	At least one ambulance should be equipped per 30,000 population (counted by area) to obtain the shortest radius and fastest response time [25].
Public Safety	Security staff shall visit at least once per week in public schools [26].

Key elements: temporal, spatial, aggregation, entity, condition, comparison

Table II: Key elements of city requirements and statistical results from 1000 real city requirements

Element	Form	Number	Example
Temporal	Dynamic Deadline	77	limit ... to one minute
	Static Deadline	98	at least once a week
	Interval	168	from 8am to 10am; within 24 hours;
	Default	657	The noise (always) should not exceed 50dB.
Spatial	PoIs/Tags	801	school area; all parks;
	Distance	650	Nearby
	Default	154	(everywhere) ; (all) locations
Aggregation	Count, Sum	256	in total; x out of N locations; %;
	Average	196	per m <sup>2</sup> ;
	Max, Min	67	highest/lowest value
Entity	Subject	1000	air quality; Buses;
Comparison	Value comparison	836	More than, less than
	Boolean	388	Street is blocked; should
	Not	456	It is unlawful/prohibited...
Condition	Until	24	keep... until the street is not blocked.
	If/Except	44	If rainy, the speed limit...

quantitatively specified city requirements (e.g., standards, regulations, city codes, and laws) across different application domains, including energy, environment, transportation, emergency, and public safety from over 70 cities (e.g. New York City, San Francisco, Chicago, Washington D.C., Beijing, etc.) around the world. Some examples of these city requirements are highlighted in Table I. We identify key required features to have in a specification language and its associated use in a city runtime monitor. The summarized statistical results of the study and key elements we identified (i.e., temporal, spatial, aggregation, entity, comparison, and condition) are shown in Table II.

**Temporal:** Most of the requirements include a variety of temporal constraints, e.g. a static deadline, a dynamic deadline, or time intervals. In many cases (65.7%), the temporal information is not explicitly written in the requirement, which usually means it should be “always” satisfied. In addition, city requirements are highly real-time driven. In over 80% requirements, cities are required to detect requirement violations at runtime. It indicates a high demand for runtime monitoring.

**Spatial:** A requirement usually specifies its spatial range explicitly using the Points of Interest (PoIs) (80.1%), such

as “park”, “xx school”, along with a distance range (65%). One requirement usually points to a set of places (e.g. all the schools). Therefore, it is very important for a formal language to be able to specify the spatial elements across many locations within the formula, rather than one formula for each location.

We also found that the city requirements specify a very large spatial scale. Different from the requirements of many other CPS, requirements from smart cities are highly spatial-specific and usually involve a very large number of locations/sensors. For example, the first requirement in Table I specifies a vehicle idling time “adjacent to any school, pre-K to 12th grade in the City of New York”. There are about 2000 pre-K to 12th schools, even counting 20 street segments nearby each school, there are 40,000 data streams to be monitored synchronously. An efficient monitoring is highly demanded.

**Aggregation:** In 51.9% cases, requirements are specified on the aggregated signal over an area, such as, “the total amount”, “average...per 100 square feet”, “up to four vending vehicles in any given city block”, “at least 20% of travelers from all entrances should ...”, etc. Therefore, aggregation is a key feature for the specification language.

#### IV. FORMALIZING TEMPORAL-SPATIAL REQUIREMENTS

SaSTL extends STL with two spatial operators: a *spatial aggregation* operator and a *neighborhood counting* operator. Spatial aggregation enables combining (according to a chosen operation) measurements of the same type (e.g., environmental temperature), but taken from different locations. The use of this operator can be suitable in requirements where it is necessary to evaluate the average, best or worst value of a signal measurement in an area close to the desired location. The neighborhood counting operator allows measuring the number/percentage of neighbors of a location that satisfy a certain requirement.

##### A. SaSTL Syntax

We define a multi-dimensional *spatial-temporal signal* as  $\omega : \mathbb{T} \times L \rightarrow \{\mathbb{R} \cup \{\perp\}\}^n$ , where  $\mathbb{T} = \mathbb{R}_{\geq 0}$ , represents the continuous time and  $L$  is the set of locations. We define  $X = \{x_1, \dots, x_n\}$

as the set of variables for each location. Each variable can assume a real value  $v \in \mathbb{R}$  or is undefined for a particular location ( $x_i = \perp$ ). We denote by  $\pi_{x_i}(\omega)$  as the projection of  $\omega$  on its component variable  $x_i \in X$ . We define  $P = \{p_1, \dots, p_m\}$  a set of propositions (e.g.  $\{\text{School, Street, Hospital, } \dots\}$ ) and  $\mathcal{L}$  a labeling function  $\mathcal{L} : L \rightarrow 2^P$  that assigns for each location the set of the propositions that are true in that location.

A *weighted undirected graph* is a tuple  $G = (L, E, \eta)$  where  $L$  is a finite non-empty set of nodes representing locations,  $E \subseteq L \times L$  is the set of edges connecting nodes, and  $\eta : E \rightarrow \mathbb{R}_{\geq 0}$  is a cost function over edges. We define the *weighted distance* between two locations  $l, l' \in L$  as

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

Then we define the spatial domain  $\mathcal{D}$  as,

$$\begin{aligned} \mathcal{D} &:= ([d_1, d_2], \psi) \\ \psi &:= \top \mid p \mid \neg \psi \mid \psi \vee \psi \end{aligned}$$

where  $[d_1, d_2]$  defines a spatial interval with  $d_1 < d_2$  and  $d_1, d_2 \in \mathbb{R}$ , and  $\psi$  specifies the property over the set of propositions that must hold in each location. Intuitively, it draws two circles with radius  $r_1 = d_1$  and  $r_2 = d_2$ , and the locations  $l \models \psi$  between these two circles are selected. In particular,  $\mathcal{D} = ([0, +\infty), \top)$  indicates the whole spatial domain. We denote  $L_{([d_1, d_2], \psi)}^l := \{l' \in L \mid 0 \leq d_1 \leq d(l, l') \leq d_2 \text{ and } \mathcal{L}(l') \models \psi\}$  as the set of locations at a distance between  $d_1$  and  $d_2$  from  $l$  for which  $\mathcal{L}(l')$  satisfies  $\psi$ . We denote the set of non-null values for signal variable  $x$  at time point  $t$  location  $l$  over locations in  $L_{\mathcal{D}}^l$  by

$$\alpha_{\mathcal{D}}^x(\omega, t, l) := \{\pi_x(\omega)[t, l'] \mid l' \in L_{\mathcal{D}}^l \text{ and } \pi_x(\omega)[t, l'] \neq \perp\}.$$

We define a set of operations  $\text{op}(\alpha_{\mathcal{D}}^x(\omega, t, l))$  for  $\text{op} \in \{\max, \min, \text{sum}, \text{avg}\}$  when  $\alpha_{\mathcal{D}}^x(\omega, t, l) \neq \emptyset$  that computes the maximum, minimum, summation and average of values in the set  $\alpha_{\mathcal{D}}^x(\omega, t, l)$ , respectively. To be noted, Graph  $G$  and its weights between nodes are constructed flexibly based on the property of the system. For example, we can build a graph with fully connected sensor nodes and their Euclidean distance as the weights when monitoring the air quality in a city; or we can also build a graph that only connects the street nodes when the two streets are contiguous and apply Manhattan distance. It does not affect the syntax and semantics of SaSTL.

The syntax of SaSTL is given by

$$\varphi := x \sim c \mid \neg \varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \mathcal{U}_I \varphi_2 \mid \mathcal{A}_{\mathcal{D}}^{\text{op}} x \sim c \mid \mathcal{C}_{\mathcal{D}}^{\text{op}} \varphi \sim c$$

where  $x \in X$ ,  $\sim \in \{<, \leq\}$ ,  $c \in \mathbb{R}$  is a constant,  $I \subseteq \mathbb{R}_{>0}$  is a real positive dense time interval,  $\mathcal{U}_I$  is the *bounded until* temporal operators from STL. The *always* (denoted  $\square$ ) and *eventually* (denoted  $\diamond$ ) temporal operators can be derived the same way as in STL, where  $\diamond \varphi \equiv \text{true } \mathcal{U}_I \varphi$ , and  $\square \varphi \equiv \neg \diamond \neg \varphi$ .

In SaSTL, we define a set of spatial aggregation operators  $\mathcal{A}_{\mathcal{D}}^{\text{op}} x \sim c$  for  $\text{op} \in \{\max, \min, \text{sum}, \text{avg}\}$  that evaluate the aggregated product of traces  $\text{op}(\alpha_{\mathcal{D}}^x(\omega, t, l))$  over a set of locations  $l \in L_{\mathcal{D}}^l$ . We also define a set of new spatial counting operators  $\mathcal{C}_{\mathcal{D}}^{\text{op}} \varphi \sim c$  for  $\text{op} \in \{\max, \min, \text{sum}, \text{avg}\}$  that counts the satisfaction of traces over a set of locations. More precisely, we define  $\mathcal{C}_{\mathcal{D}}^{\text{op}} \varphi = \text{op}(\{g((\omega, t, l') \models \varphi) \mid l' \in L_{\mathcal{D}}^l\})$ ,

where  $g((\omega, t, l) \models \varphi) = 1$  if  $(\omega, t, l) \models \varphi$ , otherwise  $g((\omega, t, l) \models \varphi) = 0$ . From the new counting operators, we also derive the *everywhere* operator as  $\boxplus_{\mathcal{D}} \varphi \equiv \mathcal{C}_{\mathcal{D}}^{\min} \varphi > 0$ , and *somewhere* operator as  $\boxtimes_{\mathcal{D}} \varphi \equiv \mathcal{C}_{\mathcal{D}}^{\max} \varphi > 0$ . In addition,  $\mathcal{C}_{\mathcal{D}}^{\text{sum}} \varphi$  specifies the total number of locations that satisfy  $\varphi$  and  $\mathcal{C}_{\mathcal{D}}^{\text{avg}} \varphi$  specifies the percentage of locations satisfying  $\varphi$ .

We now illustrate how to use SaSTL to specify various city requirements, especially for the spatial aggregation and spatial counting, and how important these operators are for the smart city requirements using examples below.

**Example 1 (Spatial Aggregation).** Assume we have a requirement, “The average noise level in the school area (within 1 km) in New York City should always be less than 50 dB and the worst should be less than 80 dB in the next 3 hours” is formalized as,  $\boxplus_{([0, +\infty), \text{School}]} \square_{[0, 3]} ((\mathcal{A}_{([0, 1], \top)}^{\text{avg}})^{x_{\text{Noise}}} < 50) \wedge (\mathcal{A}_{([0, 1], \top)}^{\max})^{x_{\text{Noise}}} < 80)$ .  $([0, +\infty), \text{School})$  selects all the locations labeled as “school” within the whole New York city  $([0, +\infty))$  (predefined by users).  $\square_{[0, 3]}$  indicates this requirement is valid for the next three hours.  $(\mathcal{A}_{([0, 1], \top)}^{\text{avg}})^{x_{\text{Noise}}} < 50) \wedge (\mathcal{A}_{([0, 1], \top)}^{\max})^{x_{\text{Noise}}} < 80$  calculates the average and maximal values in 1 km for each “school”, and compares them with the requirements, i.e. 50 dB and 80 dB.

Without the spatial aggregation operators, STL and its extended languages cannot specify this requirement. First, they are not able to first dynamically find all the locations labeled as “school”. To monitor the same spatial range, users have manually get all traces from schools, and then repeatedly apply this requirement to each located sensor within 1 km of a school and do the same for all schools. More importantly, STL and its extended languages could not specify “average” or “worst” noise level. Instead, it only monitors each single value, which is prone to noises and outliers and thereby causes inaccurate results.

**Example 2 (Spatial Counting).** A requirement that “At least 90% of the streets, the particulate matter (PM<sub>x</sub>) emission should not exceed Moderate in 2 hours” is formalized as  $\mathcal{C}_{([0, +\infty), \text{Street}]}^{\text{avg}} (\square_{[0, 2]} (x_{\text{PM}_x} < \text{Moderate})) > 0.9$ .  $\mathcal{C}_{([0, +\infty), \text{Street}]}^{\text{avg}} \varphi > 0.9$  represents the percentage of satisfaction is larger than 90%. Specifying the percentage of satisfaction is very common and important among city requirements.

## B. SaSTL Semantics

We define the SaSTL semantics as the *satisfiability relation*  $(\omega, t, l) \models \varphi$ , indicating that the spatio-temporal signal  $\omega$  satisfies a formula  $\varphi$  at the time point  $t$  in location  $l$  when  $\pi_v(\omega)[t, l] \neq \perp$  and  $\alpha_{\mathcal{D}}^x(\omega, t, l) \neq \emptyset$ . We define that  $(\omega, t, l) \models \varphi$  if  $\pi_v(\omega)[t, l] = \perp$ .

$$(\omega, t, l) \models x \sim c \quad \Leftrightarrow \pi_x(\omega)[t, l] \sim c$$

$$(\omega, t, l) \models \neg \varphi \quad \Leftrightarrow (\omega, t, l) \not\models \varphi$$

$$(\omega, t, l) \models \varphi_1 \wedge \varphi_2 \quad \Leftrightarrow (\omega, t, l) \models \varphi_1 \text{ and } (\omega, t, l) \models \varphi_2$$

$$\begin{aligned} (\omega, t, l) \models \varphi_1 \mathcal{U}_I \varphi_2 &\Leftrightarrow \exists t' \in (t + I) \cap \mathbb{T} : (\omega, t', l) \models \varphi_2 \\ &\text{and } \forall t'' \in (t, t'), (\omega, t'', l) \models \varphi_1 \end{aligned}$$

$$(\omega, t, l) \models \mathcal{A}_{\mathcal{D}}^{\text{op}} x \sim c \Leftrightarrow \text{op}(\alpha_{\mathcal{D}}^x(\omega, t, l)) \sim c$$

$$(\omega, t, l) \models \mathcal{C}_{\mathcal{D}}^{\text{op}} \varphi \sim c \Leftrightarrow \text{op}(\{g((\omega, t, l') \models \varphi) \mid l' \in L_{\mathcal{D}}^l\}) \sim c$$

where, for counting operator  $(\omega, t, l) \models \mathcal{C}_D^{\text{op}} \varphi \sim c$ , the valid ranges for  $c$  are  $c \in [0, 1]$  when  $\text{op} = \text{sum}/\text{min}$ , and  $c \in [0, N]$  when  $\text{op} = \text{sum}/\text{min}$ . Otherwise (e.g.,  $c < 0$ ), the requirement is trivially satisfied or violated.

**Example 3.** Following Example 1, checking the city states with a requirement,

$$\boxplus_{([0, +\infty), \text{Hospital})} \boxminus_{[0, 5]} ((\mathcal{A}_{([0, 500], \tau)}^{\text{avg}} x_{\text{AQI}} < 50) \wedge (\mathcal{A}_{([0, 500], \tau)}^{\text{max}} x_{\text{AQI}} < 80)),$$

to start with, assuming we have the AQI level data from a number of sensors within 500 meters of one of the hospital, the sensor readings in 5 hours as,  $\{[51, \dots, 11], [80, \dots, 30], \dots, [40, \dots, 30]\}$ ,  $\varphi_t = (\mathcal{A}_{([0, 500], \tau)}^{\text{avg}} x_{\text{AQI}} < 50) \wedge (\mathcal{A}_{([0, 500], \tau)}^{\text{max}} x_{\text{AQI}} < 80)$ , then, we check  $\varphi_t$  for this hospital at each time,

at  $t = 1$ ,  $\text{avg}(51, \dots, 40) > 50 \wedge \text{max}(51, \dots, 40) < 80$ , thus,  $\varphi_{t1} = \text{False}$ ,

at  $t = 2$ ,  $\text{avg}(49, \dots, 20) < 50 \wedge \text{max}(49, \dots, 20) > 80$ , thus,  $\varphi_{t1} = \text{False}$ ,

...

at  $t = 5$ ,  $\text{avg}(11, \dots, 30) < 50 \wedge \text{max}(11, \dots, 30) < 80$ , thus,  $\varphi_{t1} = \text{True}$ .

Thus, we have  $\boxminus_{[0, 5]} \varphi_t = \text{False}$ .

Next, the monitor checks all qualified hospitals the same way and reaches the final results,

$$\boxplus_{([0, +\infty), \text{Hospital})} \boxminus_{[0, 5]} ((\mathcal{A}_{([0, 500], \tau)}^{\text{avg}} x_{\text{AQI}} < 50) \wedge (\mathcal{A}_{([0, 500], \tau)}^{\text{max}} x_{\text{AQI}} < 80)) = \text{False}.$$

In a real scenario, the monitor algorithm can also decide to terminate the monitor and return the False result when at  $t = 1$ , because the always operator returns False as long as a one-time violation occurs. Similarly, the everywhere operator will also return False when the first hospital violates the requirement.

**Definition 1** (Quantitative Semantics). Let  $x > c$  be a numerical predicate, we then define the robustness degree (i.e. the quantitative satisfaction) function  $\rho(\varphi, \omega, t, l)$  for an SaSTL formula over a spatial-temporal signal  $\omega$  as,

$$\begin{aligned} \rho(x \sim c, \omega, t, l) &= \pi_x(\omega)[t, l] - c \\ \rho(-\varphi, \omega, t, l) &= -\rho(\varphi, \omega, t, l) \\ \rho(\varphi_1 \vee \varphi_2, \omega, t, l) &= \max\{\rho(\varphi_1, \omega, t, l), \rho(\varphi_2, \omega, t, l)\} \\ \rho(\varphi_1 \mathcal{U}_I \varphi_2, \omega, t, l) &= \sup_{t' \in (t, t+I) \cap \mathbb{T}} (\min\{\rho(\varphi_2, \omega, t', l), \\ &\quad \inf_{t'' \in [t, t']} (\rho(\varphi_1, \omega, t'', l))\}) \end{aligned}$$

$$\rho(\mathcal{A}_D^{\text{op}} x \sim c, \omega, t, l) = \begin{cases} \frac{\text{sum}(\alpha_D^x(\omega, t, l)) - c}{|\alpha_D^x(\omega, t, l)|} & \text{op} = \text{sum} \\ \text{op}(\alpha_D^x(\omega, t, l)) - c & \text{op} \in \{\text{max}, \text{min}, \text{avg}\} \end{cases}$$

$$= \begin{cases} \max_{l' \in L_D^l} \{\rho(\varphi, \omega, t, l')\} & \text{op} = \text{max} \\ \min_{l' \in L_D^l} \{\rho(\varphi, \omega, t, l')\} & \text{op} = \text{min} \\ \delta(\lceil c \rceil, \{\rho(\varphi, \omega, t, l') \mid l' \in L_D^l\}) & \text{op} = \text{sum} \\ \delta(\lceil c \times |L_D^l| \rceil, \{\rho(\varphi, \omega, t, l') \mid l' \in L_D^l\}) & \text{op} = \text{avg} \end{cases}$$

where we define  $\delta(k, S)$  as a function that returns the  $k$ th smallest number of set  $S$ ,  $|S| > 0$ , and  $0 \leq k \leq |S|$ . For  $\mathcal{C}_D^{\text{op}} \varphi \sim c$ , when  $\text{op} = \text{sum}$ , it requires that there are at least  $\lceil c \rceil$  locations that satisfy  $\varphi$ , thus, we denote the  $\lceil c \rceil$ th smallest robustness value from  $\{\rho(\varphi, \omega, t, l') \mid l' \in L_D^l\}$  as the robustness

value of this formula.  $\lceil c \rceil$  indicates the smallest integer that is larger than or equal to  $c$ . Similarly, when  $\text{op} = \text{avg}$ , the formula is converted as there are at least  $\lceil c \times |L_D^l| \rceil$  locations that satisfy  $\varphi$ , thus, we denote the  $\lceil c \times |L_D^l| \rceil$ th smallest robustness value from  $\{\rho(\varphi, \omega, t, l') \mid l' \in L_D^l\}$  as the robustness value of this formula. Same as the Boolean semantics, the valid ranges for  $c$  are  $c \in [0, 1]$  when  $\text{op} = \text{sum}/\text{min}$ , and  $c \in [0, N]$  when  $\text{op} = \text{sum}/\text{min}$ . Otherwise (e.g.,  $c < 0$ ), the requirement is trivially satisfied or violated.

**Example 4.** Assuming we have data  $(1, 2, 3)$ ,  $(2, 3, 4)$ ,  $(4, 5, 7)$  from three locations satisfying  $\mathcal{D}$ , thus,

$$\begin{aligned} \bullet \rho(\mathcal{C}_D^{\text{max}}(\boxminus_{[0, 2]}(x > 5)) > 0) &= \rho(\mathcal{C}_D^{\text{max}}(\{-4, -3, 2\}) > 0) = 2 \\ \bullet \rho(\mathcal{C}_D^{\text{min}}(\boxminus_{[0, 2]}(x > 5)) > 0) &= \rho(\mathcal{C}_D^{\text{min}}(\{-4, -3, 2\}) > 0) = -4 \\ \bullet \rho(\mathcal{C}_D^{\text{sum}}(\boxminus_{[0, 2]}(x > 5)) > 1) &= \rho(\mathcal{C}_D^{\text{sum}}(\{-4, -3, 2\}) > 1) = -3 \\ \bullet \rho(\mathcal{C}_D^{\text{avg}}(\boxminus_{[0, 2]}(x > 5)) > 0.2) &= \rho(\mathcal{C}_D^{\text{avg}}(\{-4, -3, 2\}) > 0.2) = 2 \end{aligned}$$

The quantitative semantics of SaSTL inherit the two fundamental properties of STL, i.e., soundness and correctness. We give the formal definitions below.

**Theorem 1** (Soundness). Let  $\varphi$  be an STL formula,  $\omega$  a trace and  $t$  a time,

$$\begin{aligned} \rho(\varphi, \omega, t, l) > 0 &\Rightarrow (\omega, t, l) \models \varphi \\ \rho(\varphi, \omega, t, l) < 0 &\Rightarrow (\omega, t, l) \not\models \varphi \end{aligned}$$

Secondly, if  $\omega$  satisfies  $\varphi$  at time  $t$ , any other trace  $\omega'$  whose point-wise distance from  $\omega$  is smaller than  $\rho(\varphi, \omega, t, l)$  also satisfies  $\varphi$  at time  $t$ .

**Theorem 2** (Correctness). Let  $\varphi$  be an STL formula,  $\omega$  and  $\omega'$  traces over the same time and spatial domains, and  $t, l \in \text{dom}(\varphi, \omega)$ , then

$$(\omega, t, l) \models \varphi \text{ and } \|\omega - \omega'\|_\infty < \rho(\varphi, \omega, t, l) \Rightarrow (\omega', t, l) \models \varphi$$

In summary, the qualitative value indicates if the signal (i.e. city data) satisfies the requirement. The quantitative value indicates the satisfaction or dissatisfaction degree. If it is larger or equal than zero, it means that the requirement is satisfied. The larger the value is, the more the requirement is satisfied. On the contrary, if the value is smaller than zero, it means the requirement is not satisfied. The smaller the value is, the more the requirement is dissatisfied.

## V. EFFICIENT MONITORING FOR SASTL

In this section, we first present both Boolean and quantitative monitoring algorithms for SaSTL, then describe two optimization methods to speed up the monitoring performance.

### A. Monitoring Algorithms for SaSTL

The inputs of the monitor are the SaSTL requirements  $\varphi$  (including the time  $t$  and location  $l$ ), a weighted undirected graph  $G$  and the temporal-spatial data  $\omega$ . In smart cities, the data on city states is collected continuously or periodically.

For the Boolean monitoring algorithm, the output for each requirement is a Boolean value indicating whether the requirement is satisfied or not. For the quantitative monitoring algorithm (Algorithm 1), the output for each requirement is a

**Algorithm 1:** SaSTL quantitative monitoring algorithm  
MonitorQ( $\varphi, \omega, t, l, G$ )

---

```

Input : SaSTL Requirement  $\varphi$ , Signal  $\omega$ , Time  $t$ , Location  $l$ , weighted
undirected graph  $G$ 
Output: Satisfaction Value  $\rho$ 
begin
  switch  $\varphi$  do
    Case  $x \sim c$ 
      | return  $\pi_x(\omega)[t, l] - c$ ;
    Case  $\neg\varphi$ 
      | return  $\neg$ MonitorQ( $\varphi, \omega, t, l, G$ );
    Case  $\varphi_1 \wedge \varphi_2$ 
      | return  $\min$ (MonitorQ( $\varphi_1, \omega, t, l, G$ ),
      | MonitorQ( $\varphi_2, \omega, t, l, G$ ));
    Case  $\varphi_1 \mathcal{U}_I \varphi_2$ 
      | Real  $v := -\infty$ 
      | for  $t' \in (t+I) \cap \mathbb{T}$  do
      |    $v' :=$  MonitorQ( $\varphi_2, \omega, t', l, G$ )
      |   for  $t'' \in [t, t']$  do
      |     |  $v' := \min\{v', \text{MonitorQ}(\varphi_1, \omega, t'', l, G)\}$ 
      |   end
      |    $v = \max\{v, v'\}$ 
      | end
      | return  $v$ ;
    Case  $\mathcal{A}_D^{\text{op}} x \sim c$ ;
      | return AggregateQ( $x, c, \text{op}, \mathcal{D}, t, l, G$ );
    Case  $\mathcal{C}_D^{\text{op}} \varphi \sim c$ ;
      | return CountingNeighboursQ( $\varphi, c, \text{op}, \mathcal{D}, t, l, G$ );
  end
end

```

---

number indicating the satisfaction degree of the requirement. To start with, the monitoring algorithm parses  $\varphi$  to sub-formulas and calculates the satisfaction for each operation recursively. We derived operators  $\square$  and  $\diamond$  from  $\mathcal{U}_I$ , and operators  $\boxplus$  and  $\boxtimes$  from  $\mathcal{C}_D^{\text{op}} \sim c$ , so we only show the algorithms for  $\mathcal{U}_I$  and  $\mathcal{C}_D^{\text{op}} \sim c$ .

We present the quantitative monitoring algorithms of the operators  $\mathcal{A}_D^{\text{op}}$  and  $\mathcal{C}_D^{\text{op}}$  in Algorithm 2 and Algorithm 3, respectively. We apply distributed parallel algorithm deScan() [27] to accelerate the process of searching locations that satisfy  $\mathcal{D}$ . As we can tell from the algorithms, essentially,  $\mathcal{A}_D^{\text{op}}$  calculates the aggregated values on the signal over a spatial domain, while  $\mathcal{C}_D^{\text{op}}$  calculates the aggregated results over spatial domain. For the quantitative monitoring algorithm (as presented in Algorithm 1), the *output* for each requirement is a robustness value indicating its satisfaction degree. Similar to the Boolean monitoring algorithm, the quantitative monitoring algorithm also parses  $\varphi$  to sub-formulas and calculates the satisfaction for each operation recursively.

The time complexity of monitoring the logical and temporal operators of SaSTL is the same as STL [28]. The time complexity to monitor classical logical operators or basic propositions such as  $\neg x$ ,  $\wedge$  and  $x \sim c$  is  $O(1)$ . The time complexity to monitor temporal operators such as  $\square_I$ ,  $\diamond_I$ ,  $\mathcal{U}_I$  is  $O(T)$ , where  $T$  is the total number of samples within time interval  $I$ . In this paper, we present the time complexity analysis for the spatial operators (Lemma 3) and the new SaSTL monitoring algorithm (Theorem 6). The total number of locations is denoted by  $n$ . We assume that the positions of the locations cannot change in time (a fixed grid). We can pre-compute all the distances between locations and store them in an array of range trees [29] (one range tree for each location). We further denote the monitored formula as  $\phi$ , which can be

**Algorithm 2:** AggregateQ( $x, \text{op}, \mathcal{D}, \omega, t, l, G$ )

---

```

begin
  Real  $v := 0$ ;  $n := 0$ ;
  if  $\text{op} == \text{"min"}$  then  $v := \infty$ ;
  if  $\text{op} == \text{"max"}$  then  $v := -\infty$ ;
   $L_{\mathcal{D}}^l := \text{deScan}(l, G, \mathcal{D})$ 
  for  $l' \in L_{\mathcal{D}}^l$  do
    if  $\text{op} \in \{\text{min}, \text{max}, \text{sum}\}$  then
      |  $v := \text{op}(v, \pi_x(\omega)[t, l'])$ ;
    end
    if  $\text{op} == \text{"avg"}$  then
      |  $v := \text{sum}(v, \pi_x(\omega)[t, l'])$ ;
    end
     $n := n + 1$ 
  end
  if  $n == 0$  then return  $\infty$ ;
  if  $\text{op} == \text{"avg"} \wedge n \neq 0$  then return  $v/n - c$ ;
  if  $\text{op} == \text{"sum"} \wedge n \neq 0$  then return  $(v - c)/n$ ;
  else return  $v - c$ ;
end

```

---

**Algorithm 3:** CountingNeighboursQ( $x, \text{op}, \mathcal{D}, \omega, t, l, G$ )

---

```

begin
  Real  $n := 0$ , List  $s := \text{Null}$ ;
   $L_{\mathcal{D}}^l := \text{deScan}(l, G, \mathcal{D})$ 
  for  $l' \in L_{\mathcal{D}}^l$  do
    |  $s.\text{add}(\text{Monitor}(\varphi, \omega, t, l', G))$ 
    |  $n := n + 1$ 
  end
  if  $n == 0$  then return  $\infty$ ;
  else
    switch  $\text{op}$  do
      Case  $\text{max}$ 
        | return  $s.\text{max}()$ 
      Case  $\text{min}$ 
        | return  $s.\text{min}()$ 
      Case  $\text{sum}$ 
        | return  $s.\text{max}(\text{round}(c))$ 
      Case  $\text{avg}$ 
        | return  $s.\text{max}(\text{round}(c \times n))$ 
    end
  end
end

```

---

represented by a syntax tree, and let  $|\phi|$  denote the total number of nodes in the syntax tree (number of operators).

**Lemma 1** (Complexity of spatial operators). *The time complexity to monitor at each location  $l$  at time  $t$  the satisfaction of a spatial operator such as  $\boxplus_{\mathcal{D}}$ ,  $\boxtimes_{\mathcal{D}}$ ,  $\mathcal{A}_{\mathcal{D}}^{\text{op}}$ , and  $\mathcal{C}_{\mathcal{D}}^{\text{op}}$  is  $O(\log(n) + |L|)$  where  $L$  is the set of locations at distance within the range  $\mathcal{D}$  from  $l$ .*

**Theorem 3.** *The time complexity of the SaSTL monitoring algorithm is upper-bounded by  $O(|\phi| \times T_{\text{max}} \times (\log(n) + |L|_{\text{max}}))$  where  $T_{\text{max}}$  is the largest number of samples of the intervals considered in the temporal operators of  $\phi$  and  $|L|_{\text{max}}$  is the maximum number of locations defined by the spatial temporal operators of  $\phi$ .*

**B. Performance Improvement of SaSTL Parsing**

To monitor a requirement, the first step is parsing the requirement to a set of sub formulas with their corresponding spatial-temporal ranges. Then, we calculate the results for the sub-formulas. The traditional parsing process of STL builds and calculates the syntax tree on the sequential order of the formula. It does not consider the complexity of each sub-formula. However, in many cases, especially with the PoIs specified in smart cities, checking the simpler propositional variable to quantify the spatial domain first can significantly




Figure 2: An example of city abstracted graph. A requirement is  $\boxplus_{([0,+\infty),\text{School}]} \square_{[a,b]} (\mathcal{A}_{([0,d],\tau)}^{\text{op}} \varphi \sim c)$  (The large nodes represent the locations of PoIs, among which the red ones represent the schools, and blue ones represent other PoIs. The small black nodes represent the locations of data sources.)

reduce the number of temporal signals to check in a complicated formula. For example, the city abstracted graph in Figure 2, the large nodes represent the locations of PoIs, among which the red ones represent the schools, and blue ones represent other PoIs. The small black nodes represent the locations of data sources (e.g. sensors). Assuming a requirement  $\boxplus_{([0,+\infty),\text{School}]} \square_{[a,b]} (\mathcal{A}_{([0,d],\tau)}^{\text{op}} \varphi \sim c)$  requires to aggregate and check  $\varphi$  only nearby schools (i.e., the red circles), but it will actually check data sources of all nearby 12 nodes if one is following the traditional parsing algorithm. In New York City, there are about 2000 primary schools, but hundreds of thousands of PoIs in total. A very large amount of computing time would be wasted in this way.

---

**Algorithm 4:** Satisfaction of  $(\varphi_1 \wedge \varphi_2, \omega)$ 


---

```

case  $\varphi_1 \wedge \varphi_2$  do
  return Monitor( $\varphi_1, \omega, t, l, G$ )  $\wedge$  Monitor( $\varphi_2, \omega, t, l, G$ );
  if  $\text{cost}(\varphi_1, l, G) \leq \text{cost}(\varphi_2, l, G)$  then
    if  $\neg \text{Monitor}(\varphi_1, \omega, t, l, G)$  then
      return Monitor( $\varphi_2, \omega, t, l, G$ );
    end
    return True;
  end
  if  $\neg \text{Monitor}(\varphi_2, \omega, t, l, G)$  then
    return Monitor( $\varphi_1, \omega, t, l, G$ );
  end
  return True;
end

```

---

To deal with this problem, we now introduce a monitoring cost function  $\text{cost} : \Phi \times L \times G_L \rightarrow \mathbb{R}^+$ , where  $\Phi$  is the set of all the possible SaSTL formulas,  $L$  is the set of locations,  $G_L$  is the set of all the possible undirected graphs with  $L$  locations. The cost function for  $\varphi$  is defined as:

$$\text{cost}(\varphi, l, G) = \begin{cases} 1 & \text{if } \varphi := p \vee \varphi := x \sim c \vee \varphi := \text{True} \\ 1 + \text{cost}(\varphi_1, l, G) & \text{if } \varphi := \neg \varphi_1 \\ \text{cost}(\varphi_1, l, G) + \text{cost}(\varphi_2, l, G) & \text{if } \varphi := \varphi_1 * \varphi_2, * \in \{\wedge, \cup\} \\ |L_{\mathcal{D}}^l| & \text{if } \varphi := \mathcal{A}_{\mathcal{D}}^{\text{op}} x \sim c \\ |L_{\mathcal{D}}^l| \text{cost}(\varphi_1, l, G) & \text{if } \varphi := \mathcal{C}_{\mathcal{D}}^{\text{op}} \varphi_1 \sim c \end{cases}$$

Using the above function, the cost of each operation is calculated before “switch  $\varphi$ ” (refer to Algorithm 1). The cost function measures how complex it is to monitor a particular SaSTL formula. This can be used when the algorithm evaluates the  $\wedge$  operator and it establishes the order in which the sub-formulas should be evaluated. The simpler sub-formula is the

first to be monitored, while the more complex one is monitored only when the other sub-formula is satisfied. We update  $\text{monitor}(\varphi_1 \wedge \varphi_2, \omega)$  in Algorithm 4. With this cost function, the time complexity of the monitoring algorithm is reduced to  $O(|\phi| \times T_{\max} \times (\log(n) + |L'_{\max}|))$ , where  $|L'|$  is the maximal number of locations that an operation is executed with the improved parsing method. The improvement is significant for city requirements, where  $|L'_{\max}| < 100 \times |L_{\max}|$ .

### C. Parallelization

In traditional STL monitor algorithm, the signals are checked sequentially. For example, to see if the data streams from all locations satisfy  $\boxplus_{\mathcal{D}} \square_{[a,b]} \varphi$  in Figure 2, usually, it would first check the signal from location 1 with  $\square_{[a,b]} \varphi$ , then location 2, and so on. At last, it calculates the result from all locations with  $\boxplus_{\mathcal{D}}$ . In this example, checking all locations sequentially is the most time-consuming part, and it could reach over 100 locations in the field.

To reduce the computing time, we parallelize the monitoring algorithm in the spatial domain. To briefly explain the idea: instead of calculating a sub-formula  $(\square_{[a,b]} \varphi)$  at all locations sequentially, we distribute the tasks of monitoring independent locations to different threads and check them in parallel. (Algorithm 5 presents the parallel version of the spatial counting operator  $\mathcal{C}_{\mathcal{D}}$ .) To start with, all satisfied locations  $l' \in L_{\mathcal{D}}^l$  are added to a task pool (a queue). In the mapping process, each thread retrieves monitoring tasks (i.e., for  $l_i, \square_{[a,b]} \varphi$ ) from the queue and executes them in parallel. All threads only execute one task at one time and is assigned a new one from the pool when it finishes the last one, until all tasks are executed. Each task obtains the satisfaction of  $\text{Monitor}(\varphi, \omega, t, l, G)$  function, and calculates the local result  $v_i$  of operation  $\text{op}()$ . The reduce step sums all the parallel results and calculates a final result of  $\text{op}()$ .

---

**Algorithm 5:** Parallelization of Counting of  $(x, \text{op}, \mathcal{D}, \omega, t, l, G)$ 


---

<pre> Function CountingNeighbours(<math>\varphi, \text{op}, \mathcal{D}, \omega, t, l, G</math>) begin   begin     paratasks = Queue();     for <math>l' \in L_{\mathcal{D}}^l</math> do       paratasks.add(<math>l'</math>);     end     results = Queue();     for <math>i</math> in <math>1..NumThreads</math> do       Thread<math>_i</math> <math>\leftarrow</math>         worker(<math>\varphi, \omega, t, l, G</math>);     end     Wait();     return op(results);   end </pre>	<pre> Function worker(<math>\varphi, \omega, t, l, G</math>):   Real <math>v := 0</math>;   if <math>\text{op} == \text{"min"}</math> then <math>v := \infty</math>;   ;   if <math>\text{op} == \text{"max"}</math> then     <math>v := -\infty</math>; ;   while <math>Num(tasks) &gt; 0</math> do     <math>l = \text{paratasks.pop}()</math>;     moni =       Monitor(<math>\varphi, \omega, t, l, G</math>);     <math>v = \text{op}(v, \text{moni})</math>;   end   results.add(<math>v</math>) </pre>
---	---

---

**Lemma 2.** *The time complexity of the parallelized algorithm  $\text{Monitor}(\phi, \omega)$  is upper bounded by  $O(|\phi| T_{\max} (\log(n) + \frac{|L_{\max}|}{P}))$  when distributed to  $P$  threads.*

In general, the parallel monitor on the spatial domain reduces the computational time significantly. It is very helpful to support runtime monitoring and decision making, especially for a large number of requirements to be monitored in a short time. In practice, the computing time also depends

on the complexity of temporal and spatial domains as well as the amount of data to be monitored. A comprehensive experimental analysis of the time complexity is presented in Section VII.

## VI. TOOL FOR THE SASTL MONITOR

We develop a user-friendly prototype tool for the SaSTL monitor that can support decision making of different stakeholders in smart cities. The interface and flowchart of the tool are shown in Figure 3. The tool allows users (e.g., city decision maker, citizens) without any formal method background to check the city performance (data) with their own requirements easily in four steps.

*Step 1: selecting the monitoring city and PoI.* To start with, users select the areas (such as a city, or a particular area of the city) to monitor, then choose the important labels that a requirement is involved with, such as, schools, parks, theaters, etc. Once selected, the important points of interest (PoIs) are shown on the map. This helps users define and verify the monitoring locations. If a location or label is not included, users are also able to add them with their GPS coordinates. The map displays the locations of the specified labels and sensors. Users can enlarge the map to check the distribution of sensors and PoIs and revise the requirements accordingly.

*Step 2: setting up the city data interface.* The data of the city states collected from sensors across temporal and spatial domains are introduced to the monitor in the Data section. For the offline monitoring, users can specify the data location of each variable on the computer. For runtime monitoring, the sensing data continuously come into the computer, the data interface of which can be set up in this section.

*Step 3: specifying the city safety requirements.* As the next important step, users specify all requirements in the requirement section. Users first select the template and then choose/fill in the essential part using the structured template language. To be noted, the entities and spatial ranges correspond to the available data variables and PoIs inputs from the areas and data sections.

We define a series of templates using structured language learning from the existing city requirements, as shown in Figure 4. The goal of these templates is to help and inspire users to specify requirements precisely. These templates are adequate to represent all the example requirements given in Table I as well as the total set of 1,000 quantitatively-defined requirements. We define the templates in a recursive way.  $T$  is a template, and  $T_1$  and  $T_2$  are instances of  $T$ . The elements in  $T$  are optional, i.e.  $\langle \rangle$  can be defined as blank, indicating this element is not applicable or default in this requirement. For example, an environmental requirement is written as, "The  $\langle$ average $\rangle$   $\langle$ air quality $\rangle$  within  $\langle$ 1 $\rangle$  mile of all  $\langle$ parks $\rangle$  should  $\langle$ always $\rangle$  be  $\langle$ above $\rangle$   $\langle$ good $\rangle$ ." The duration is interpreted as always (default) and there is no condition element. To convert a structured requirement to SaSTL, we extract the pre-defined key elements and translate them to the SaSTL formula following the rules. Meanwhile, users are also able to use the advanced features to input the city requirements in the format of the SaSTL formal formulas directly.

*Step 4: runtime monitoring.* With all the data and requirements well defined, users can start the monitor in order to check if the incoming data from the smart city satisfies the requirements. The results are displayed with a Boolean value indicating if the requirement is satisfied and a robustness value indicating how much the requirement is satisfied or violated. In addition, the map also displays the monitor results visually. Two examples are shown in Figure 5. The first one is monitoring an air quality requirements of high schools in Chicago, and the second one is monitoring a traffic requirement in New York City. The green circle represents the location satisfied the requirement and the red circle represents the location violates the requirement; the size of the circle represents the degree of satisfaction or violation. Users can zoom in and out the map to focus on a specific area or check the overall performance as needed (See Figure 5 (2)).

In summary, we defined templates helping users to specify requirements to the SaSTL formal formulae. We believe these templates can not only help users to convert the requirement from English to formal formulae, they are also helpful for users to write the requirements much more specifically and precisely. The templates defined in this paper are not sufficient to cover all the city requirements, especially the new requirements coming with more and more smart services being developed. However, the approach that using structured language to specify requirements proposed in this paper is general and effective. Also, the templates are easily extended to adapt to new requirements.

We envision this tool can be used by different stakeholders in smart cities, including but not limited to,

*City managers and decision makers:* In the city operating center, with city data collected in real time, the Tool is able to help city managers and decision makers to monitor the data at runtime. It also helps the city center to detect conflicts, and provide support for decision makers by showing the trade-offs of satisfaction degrees among potential solutions.

*City planners:* City planners, either from the government to make long-term policies or from a company to make a short-term event plan, they are able to use the Tool to verify the past city data with their requirements and make plans to prevent the violations.

*Service designers:* Smart services are designed by different stakeholders including the government, companies and private parties, they are not aware of all the other services. However, with the monitor, they can test the influence of their services on the city and adjust the services to better serve the city.

*Everyday citizens:* The tool can also provide a service to the everyday citizens. Citizens without any technical background are able to specify their own requirements and check them with the city data to find out in which areas of the city and period of the day their requirements are satisfied, and make plans about their daily life. For example, a citizen can specify an environmental requirement with his/her preferred air quality index and traffic conditions, and check the city data with the requirements and make up travelling agenda accordingly.



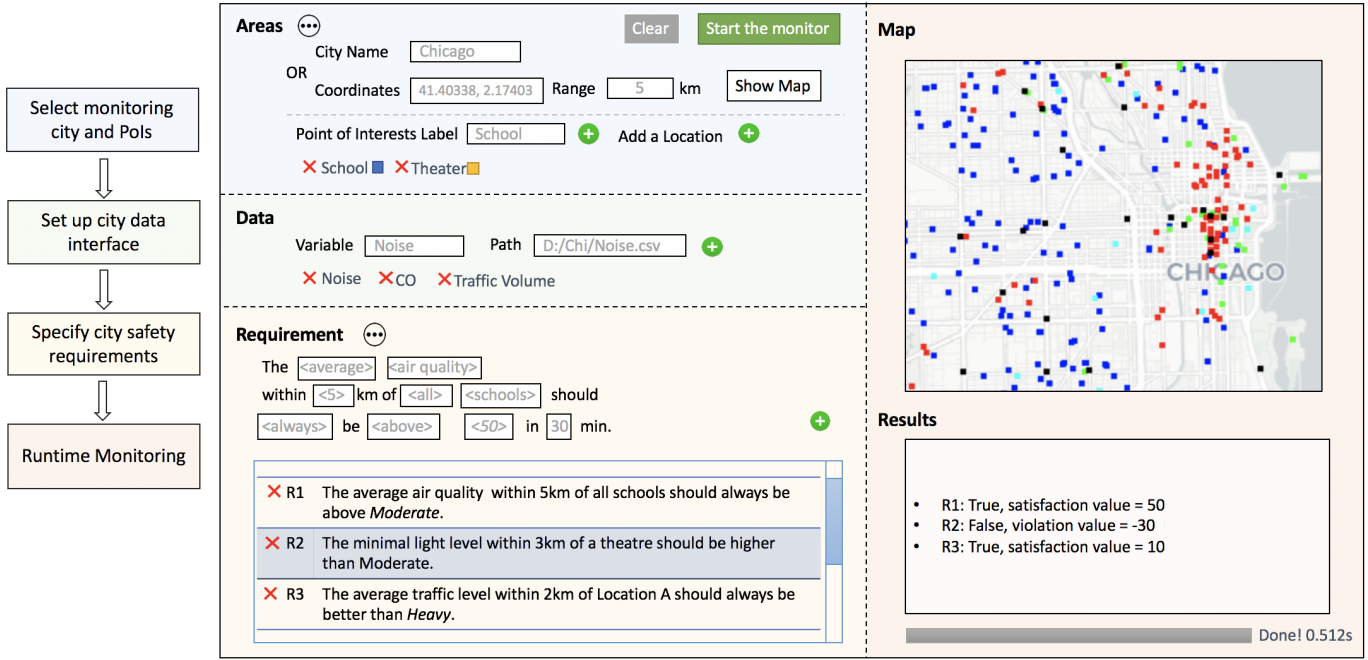


Figure 3: Interface of the SaSTL monitoring tool

T:= The <aggregation operator> <entity> within <d> miles (from <a>th mile to <b>th mile) of <spatial operator> <PoIs> should <temporal operator> be <compare> <parameter> within <t> hours (from <m>th hour to <n>th hour / on <date> day).

T:= If T1, then T2.

T:= It is prohibited that T1.

T:= T1 and/until/except T2.

Figure 4: Templates to specify city requirements

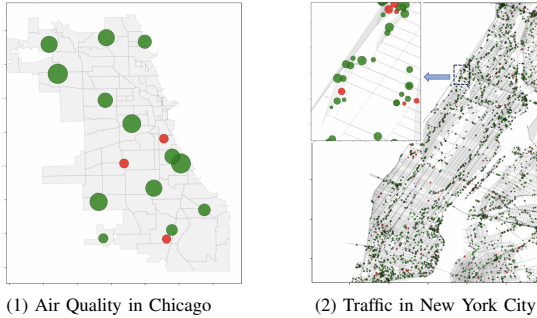


Figure 5: Display of the Monitoring Results on the Maps (The green circle represents the location satisfied the requirement and the red circle represents the location violates the requirement; the size of the circle represents the degree of satisfaction or violation.)

## VII. EVALUATION

We evaluate the SaSTL monitor by applying it to three big city application scenarios, *New York*, *Chicago*, and *Aarhus*. The experiments are evaluated on a server machine with 20 CPUs, each core is 2.2GHz, and 4 Nvidia GeForce RTX 2080Ti GPUs. The operating system is Centos 7.




Figure 6: Requirement Satisfaction Rate during Different Time Periods in Chicago

### A. Runtime Monitoring of Real-Time Requirements in Chicago

1) *Introduction:* We apply SaSTL to monitor the real-time requirements in Chicago. The framework is the same as shown in Figure 1, where we first formalize the city requirements to SaSTL formulas and then monitor the city states with the formalized requirements. Chicago is collecting and publishing city environment data (e.g., CO, NO, O3, visible light) every day since January, 2017 [1]. In our evaluation, we emulate the Chicago data as it arrives in real time, i.e. assuming the city was operating with our SaSTL monitor. Specifically, we monitor data from 118 locations between January, 2017 and May, 2019. In addition, we incorporate the Chicago crime rate data published by the city of Chicago [30]. The sampling rates of sensors vary by locations and variables (e.g., CO is updated every few seconds, and the crime rate map is updated by events), so we normalize the data frequency as one minute. Then we specify 80 safety and performance requirements that are generated from the real requirements, and apply the SaSTL to monitor the data every 3 hours continuously to identify the requirement violations.




Figure 7: Number of Requirements Checked on Different Computing Time

2) *Chicago Performance*: Valuable information is identified from the monitor results of different periods during a day. We randomly select 30 days of weekdays and 30 days of weekends. We divide the daytime of a day into 4 time periods and 3 hours per time period. We calculate the percentage of satisfaction (i.e., number of satisfied requirement days divides 30 days) for each time period, respectively. The results of two example requirements CR1 and CR2 are shown in Figure 6. CR1 specifies “The average air quality within 5km of all schools should always be above *Moderate* in the next 3 hours.” and is formalized as  $\boxplus_{([0,+\infty),\text{School}]} \boxminus_{[0,3]} (\mathcal{A}_{([0,5],\tau)}^{\text{avg}} x_{\text{air}} > \text{Moderate})$ . CR2 specifies “For the blocks with a high crime rate, the average light level within 3 km should always be *High*” and is formalized as  $\boxplus_{([0,+\infty),\tau]} \boxminus_{[0,3]} (x_{\text{Crime}} = \text{High} \rightarrow \mathcal{A}_{([0,3],\tau)}^{\text{avg}} x_{\text{Light}} \geq \text{High})$ .

The SaSTL monitor results can be potentially used by different stakeholders.

First, *with proper requirements defined, the city decision makers are able to identify the real problems and take actions to resolve or even avoid the violations in time*. For example, from the two example requirements in Figure 6, we could see over 20% of the time the requirements are missed everyday. Based on the monitoring results of requirement CR1, decision makers can take actions to redirect the traffic near schools and parks to improve the air quality. Another example of requirement CR2, the satisfaction is much higher (up to 33% higher in CR2, 8pm - 11pm) over weekends than workdays. There are more people and vehicles on the street on weekends, which as a result also increases the lighted areas. However, as shown in the figure, the city lighting in the areas with high crime rate is only 60%. An outcome of this result for city managers is that they should pay attention to the illumination of workdays or the areas without enough light to enhance public safety.

Second, *it gives the citizens the ability to learn the city conditions and map that to their own requirements*. They can make decisions on their daily living, such as the good time to visit a park. For example, requirement CR1, 11am - 2pm has the lowest satisfaction rate of the day. The instantaneous air quality seems to be fine during rush hour, but it has an accumulative result that affects citizens’ (especially students and elderly people) health. A potential suggestion for citizens who visit or exercise in the park is to avoid 11am - 2pm.

3) *Algorithm Performance*: We count the average monitoring time taken by each requirement when monitoring for 3-hour data. Then, we divide the computing time into 5 categories, i.e., less than 1 second, 1 to 10 seconds, 10 to

60 seconds, 60 to 120 seconds, and longer than 120 seconds, and count the number of requirements under each category under the conditions of standard parsing, improved parsing with single thread, 4 threads, and 8 threads. The results are shown in Figure 7. Comparing the 1st (standard parsing) and 4th (8 threads) bar, without the improved monitoring algorithms, for about 50% of the requirements, each one takes more than 2 minutes to execute. The total time of monitoring all 80 requirements is about 2 hours, which means that the city decision maker can only take actions to resolve the violation at earliest 5 hours later. However, with the improved monitoring algorithms, for 49 out of 80 requirements, each one of them is executed within 60 seconds, and each one of the rest requirements is executed within 120 seconds. The total execution time is reduced to 30 minutes, which is a reasonable time to handle as many as 80 requirements. More importantly, it illustrates the effectiveness of the parsing function and parallelization methods. Even if there are more requirements to be monitored in a real city, it is doable with our approach by increasing the number of processors.

## B. Runtime Conflict Detection and Resolution in Simulated New York City

1) *Introduction*: The framework of runtime conflict detection and resolution [31], [10] considers a scenario where smart services send action requests to the city center, and where a simulator predicts how the requested actions change the current city states over a finite future horizon of time. Then it checks the predicted states against city requirements. If the requirements are satisfied, the requested actions will be approved to execute in the city. If there exists a requirement violation within the future horizon, a conflict is detected. CityResolver will be applied to resolve the conflicts. Details of the resolution are not the main part of this paper, please refer to CityResolver [10]. Note that with the conflicts detected and resolved, the city’s future states will be affected. In this paper, we apply the SaSTL monitor to specify requirements with spatial aggregation and check the *predicted spatial-temporal data* with the SaSTL formulas.

We set up a smart city simulation of New York City using the Simulation of Urban MObility (SUMO) [32] with the traffic pattern (vehicle in-coming rate of key streets) from real city data [33], on top of which, we implement 10 services (S1: Traffic Service, S2: Emergency Service, S3: Accident Service, S4: Infrastructure Service, S5: Pedestrian Service, S6: Air Pollution Control Service, S7: PM2.5/PM10 Service, S8: Parking Service, S9: Noise Control Service, and S10: Event Service). The real-time states (including CO, NO, O3, PMx, Noise, Traffic, Pedestrian Number, Signal Lights, Emergency Vehicles, and Accident number) from the domains of environment, transportation, events and emergencies are obtained from about 10,000 simulated nodes. Then, we apply the STL Monitor as the *baseline* to compare the capability of requirement specification and the ability to improve city performance. We simulate the city running for 30 days with sampling rate as 10 seconds in two control sets, one without any monitor and one with the SaSTL monitor. For the first set (no monitor),

Table III: Safety and Performance Requirements for New York City

	Requirement	SaSTL
<b>NYR1</b>	The average noise level in the school area (within 1km) should always be less than 50dB in the next 30min.	$\Box_{([0,+\infty),\text{School}]} \Box_{[0,30]} (\mathcal{A}_{([0,1],\tau)}^{\text{avg}})^{x_{\text{Noise}}} < 50)$
<b>NYR2</b>	If an accident happens, at least one of the nearby hospitals (within 5km), its traffic condition within 2km should not reach the level of congestion in the next 60 min.	$\Box_{([0,+\infty),\tau]} (\text{Accident} \rightarrow \mathcal{C}_{([0,5],\text{Hospital})} (\Box_{[0,60]} (\mathcal{A}_{([0,2],\tau)}^{\text{avg}})^{x_{\text{Congestion}}} > 0))$
<b>NYR3</b>	If there is an event, the max number of pedestrians waiting at an intersection should not be greater than 50 for more than 10 minutes.	$\Box_{([0,+\infty),\tau]} (\text{Event} \rightarrow \Box_{[0,10]} (\mathcal{A}_{([0,1],\tau)}^{\text{max}})^{x_{\text{ped}}} < 50))$
<b>NYR4</b>	At least 90% of the streets, the PMx emission should not exceed <i>Moderate</i> in 60 min.	$\mathcal{C}_{([0,+\infty),\tau]}^{\text{avg}} (\Box_{[0,60]} (\mathcal{A}_{([0,1],\tau)}^{\text{max}})^{x_{\text{PMx}}} < \text{Moderate})) > 0.9$
<b>NYR5</b>	If an accident happens, it should be solved within 60 min, and before that nearby (500 m) traffic should be above moderate on average and safe in worst case.	$\Box_{([0,+\infty),\tau]} (\text{Accident} \rightarrow (\mathcal{A}_{([0,500],\tau)}^{\text{avg}})^{x_{\text{traffic}}} < \text{Moderate} \wedge \mathcal{A}_{([0,500],\tau)}^{\text{max}})^{x_{\text{traffic}}} < \text{Safe})) \wedge \Box_{[0,60]} \neg \text{Accident}$

Table IV: Comparison of the City Performance with the STL Monitor and the SaSTL Monitor

	No Monitor	SaSTL Monitor
<b>Number of Violation</b>	Unknown	<b>173</b>
<b>Air Quality Index</b>	67.91	<b>40.18</b>
<b>Noise (db)</b>	73.32	<b>41.42</b>
<b>Emergency Waiting Time (s)</b>	20.32	<b>11.88</b>
<b>Vehicle Waiting Number</b>	22.7	<b>12.6</b>
<b>Pedestrian Waiting Time (s)</b>	190.2	<b>61.1</b>
<b>Vehicle Waiting Time (s)</b>	112.12	<b>59.22</b>

there is no requirement monitor implemented. For the second one (SaSTL monitor), five examples of different types of real-time requirements and their formalized SaSTL formulas are given in Table III.

2) *NY City Performance*: The results are shown in Table IV. We measure the city performance from the domains of transportation, environment, emergency and public safety using the following metrics, the total number of violations detected (i.e., the total number of safety requirements violated during the whole simulation time), the average CO (mg) emission per street, the average noise (dB) level per street, the emergency vehicles waiting time per vehicle per intersection, the average number and waiting time of vehicles waiting in an intersection per street, and the average pedestrian waiting time per intersection.

We make some observations by comparing and analyzing the monitoring results.

First, *the SaSTL monitor obtains a better city performance with fewer number of violations detected under the same scenario*. As shown in Table IV, on average, the framework of conflict detection and resolution with the SaSTL monitor improves the air quality by 40.8%, and improves the pedestrian waiting time by 47.2% comparing to the one without a monitor.

Second, *the SaSTL monitor reveals the real city issues, helps refine the safety requirements in real time and supports improving the design of smart services*. We also compare the number of violations on each requirement. The results (Figure 8 (1)) help the city managers to measure city's performance with smart services for different aspects, and also help policymakers to see if the requirements are too strict to be satisfied by the city and make a more realistic requirement if necessary. For example, in our 30 days simulation, apparently, NYR4 on air pollution is the one requirement that is violated by most of the smart services. Similarly, Figure 8 (2) shows the number of violations caused by different smart services. Most of the violations are caused by S1, S5, S6, S7, and S10.

Table V: Computing time of requirements with standard parsing function, with improved parsing functions and different number of threads

	Standard Parsing (s)	1 thread (s)	4 threads (s)	8 threads (s)
<b>NYR1</b>	2102.13	140.29	50.31	<b>26.12</b>
<b>NYR2</b>	55.2	<b>0.837</b>	1.023	0.912
<b>NYR3</b>	69.22	22.25	7.54	<b>4.822</b>
<b>NYR4</b>	390.19	390.19	100.23	<b>53.32</b>
<b>NYR5</b>	61.76	61.76	20.25	<b>15.68</b>
<b>Total</b>	2678.5	615.32	179.35	<b>100.85</b>




Figure 8: Distributions of the violations over requirements and smart services

The five major services in total cause 71.3% of the violations. City service developers can also learn from these statistics to adjust the requested actions, the inner logic and parameters of the functions of the services, so that they can design a more compatible service with more acceptable actions in the city.

3) *Algorithm Performance*: We compare the average computing time for each requirement under four conditions, (1) using the standard parsing algorithm without the cost function, (2) improved parsing algorithm with a single thread, (3) improved parsing algorithm with spatial parallelization using 4 threads and (4) using 8 threads. The results are shown in Table V.

First, the improved parsing algorithm reduces the computing time significantly for the requirement specified on PoIs, especially for NYR1 that computing time reduces from 2102.13 seconds to 140.29 seconds (about 15 times). Second, the parallelization over spatial operator further reduces the computing time in most of the cases. For example, for NYR1, the computing time is reduced to 26.12 seconds with 8 threads while 140.29 seconds with single thread (about 5 times). When the amount of data is very small (NYR2), the parallelization time is similar to the single thread time, but still much efficient than the standard parsing.

The results demonstrate the effectiveness and importance




Figure 9: Comparisons of Satisfaction Rate on AR1 to AR5

of the efficient monitoring algorithms. In the table, the total time of monitoring 5 requirements is reduced from 2678.5 seconds to 100.85 seconds. In the real world, when multiple requirements are monitored simultaneously, the improvement is extremely important for real-time monitoring.

### C. Evaluation for Aarhus

1) *Introduction:* In this case study, we monitor the past data of events and states from Aarhus to show how the SaSTL monitor helps to understand the effects caused by events and therefore aids in decision making for city events. We utilize 60 days (August to September 2014) of Aarhus city data collected simultaneously across the domains of transportation (e.g., traffic volume, parking), events (e.g., cultural events and library events) and the environment (generated pollution and weather). All the data were collected from 449 observation points and published by CityPulse [34]. Data was collected with different sampling rates (e.g., the traffic data were aggregated by 5 minutes and events data were recorded by the event time), thus for the monitoring purpose, we normalize the data frequency as 5 minutes. Five safety and performance requirements and their corresponding SaSTL statements are presented with a high demand for aggregations specified for Aarhus in Table VI. Basically, AR1 to AR5 specify that when there is an event, there is a different level of safety requirements on the traffic under different circumstances. For example, AR2 focuses on the areas nearby an event, AR3 focuses on the safety of school with an event, and R4 considers the effects from extreme weather conditions. AR5 has a big picture on all schools across the city when a large cultural event is happening.

2) *Performance:* The monitoring results from Aarhus are shown in Figure 9. The percentage of satisfaction equals to the number of requirement satisfied days divided by 60 days. The following are observations on the requirements and monitoring results.

- Comparing the monitoring results on AR1 and AR2, AR1 has a much lower satisfaction rate. It also leads to a higher and reliable satisfaction rate.
- Comparing to AR2, for the same events, AR3 moves its focus on the area nearby schools. The results, however, are lower than AR2. It means that events have more influence on the school areas, which should draw attention from the city managers. Students should reduce or avoid activities during this time when there is an event going on nearby.
- During 11am to 2pm, the overall performance on all five requirements are worst, even less than 50%. It is actually

the time period right after a morning event or before an afternoon event. The monitoring results help the city managers have a better view of the distribution of effects from events.

- We also find that the satisfaction rate is very high (almost 100%) after 8pm. The reasons for that are the schools are usually closed at that time, and most of cultural and library events happen during the day. In other cities or events, the distribution will be different. However, the SaSTL monitor is general enough to help citizens and managers detect it.

The evaluation on Aarhus shows how the SaSTL monitor helps the city to understand the effects on the city from events and make better plans for events. Usually, areas with an event get caught up in complicated situations, such as paralyzed traffic, long queues with a large amount of people, emergencies and accidents. Therefore, playing back and analyzing the city data during events is extremely important for cities to avoid emergency situations for future events.

## VIII. COVERAGE ANALYSIS

We compare the specification coverage on 1000 quantitatively-specified real city requirements between STL, SSTL, STREL and SaSTL. The study is conducted by graduate students following the rules that if the language is able to specify the whole requirement directly with one single formula, then it is identified as True. To be noted, another spatial STL, SpaTeL is not considered as a baseline here, because it is not applicable to most of city spatial requirements. SpaTeL is built on a quad tree, and able to specify directions rather than the distance.

STL is only able to specify 184 out of 1000 requirements, while SSTL and STREL are able to formalize 431 requirements. SaSTL is able to specify 950 out of 1000 requirements. In particular, we made the following observations from the results. First, 50 requirements cannot be specified using any of the four languages because they are defined by complex math formulas that are ambiguous with missing key elements, relevant to the operations of many variables, or referring to a set of other requirements, e.g. “follow all the requirements from Section 201.12”, etc. Secondly, SSTL, STREL and SaSTL outperformed STL in terms of requirements with spatial ranges, such as “one-mile radius around the entire facility”; Third, SSTL and STREL have the same coverage on the requirements that only contain a temporal and spatial range. Comparing to SSTL and SaSTL, STREL can also be applied to dynamic graph and check requirements reachability, which is very useful in applications like wireless sensor networks, but not common in smart city requirements; Fourth, the rest of the requirements (467 out of 1000) measure the aggregation of a set of locations, which can only be specified using SaSTL.

## IX. RELATED WORK

Monitoring spatial-temporal properties over CPS executions has been initially investigated in [35], [36], where the authors introduced a spatial-temporal event-based model for monitoring CPS. In this model, events are labeled with time

Table VI: Safety and Performance Requirements for Aarhus

	Requirement	SaSTL
AR1	If there is an event, the traffic level nearby should always be better than <i>Moderate</i> .	Event $\rightarrow \@_{\mathcal{D}} \square_{[0,3]} x_{\text{traffic}} > \text{Moderate}$
AR2	If there is an event, the average traffic level nearby should always be better than <i>Moderate</i> and the maximum traffic level nearby should be better than <i>Safe</i> .	Event $\rightarrow \@_{\mathcal{D}} \square_{[0,3]} (\mathcal{A}_{([0,1],\tau)}^{\text{avg}} x_{\text{traffic}} > \text{Moderate} \wedge \mathcal{A}_{([0,1],\tau)}^{\text{max}} x_{\text{traffic}} > \text{Safe})$
AR3	If there is an event, the average traffic near the school (3km) should always be better than <i>Moderate</i> and the maximum traffic level should be better than <i>Heavy</i> .	Event $\rightarrow \@_{\mathcal{D}} \square_{[0,3]} (\mathcal{A}_{([0,1],\tau)}^{\text{avg}} x_{\text{traffic}} > \text{Moderate} \wedge \mathcal{A}_{([0,1],\tau)}^{\text{max}} x_{\text{traffic}} > \text{Heavy}) \wedge \text{School}$
AR4	If there is an event and the weather is rainy or snowy heavily, the average traffic level around school should be better than <i>Heavy</i> .	Event $\wedge \text{Humidity} > 50\% \rightarrow \@_{\mathcal{D}} \square_{[0,3]} (\mathcal{A}_{([0,1],\tau)}^{\text{avg}} x_{\text{traffic}} > \text{Heavy}) \wedge \text{School}$
AR5	With big cultural events going on the city, over the city, 80% schools' average traffic volume nearby (3km) should always be better than <i>Moderate</i> .	Event $\rightarrow \mathcal{C}_{\mathcal{D}} \square_{[0,3]} (\mathcal{A}_{([0,1],\tau)}^{\text{avg}} x_{\text{traffic}} > \text{Moderate}) \wedge \text{School} > 80\%$

and space stamps. These events can be triggered by actions, exchange of messages or physical changes. A centralized monitor is then responsible to process all these events. Their approach provides an algorithmic framework enabling a user to develop manually a monitor, but they do not provide any spatial-temporal specification language. The literature instead offers several logic-based specification languages to reason about the spatial structure of concurrent systems [37], medical images [38], and the topological [39] or directional [40] aspects of the interacting components. However, these logics are not practical for monitoring CPS, because they are generally computationally complex [40] or even undecidable [41].

Specification-based monitoring of spatial-temporal properties over CPS executions has become practical only recently with SpaTeL [42] and SSSL [12]. SpaTeL extends the Signal Temporal Logic [11] (STL) with the Tree Spatial Superposition Logic (TSSL) [43], [44]. TSSL classifies and detects spatial patterns by reasoning over-quad trees, suitable spatial data structures that are constructed by recursively partitioning the space into uniform quadrants. The notion of superposition in TSSL [44] provides a way to describe statistically the distribution of discrete states in a particular partition of the space and the spatial operators corresponding to *zooming in and out in* a particular region of the space. By nesting these operators, it is possible to specify self-similar and fractal-like structures [45] that generally characterize the patterns emerging in nature such as the electrical spiral formation in cardiac tissues [46]. The procedure allows one to capture very complex spatial structures, but at the price of a complex formulation of spatial properties, which are in practice only learned from some template image.

SSSL [12] extends STL with several spatial operators (i.e., somewhere, everywhere, and surround). The SSSL semantics operates on a weighted undirected graph, where the weight on each edge represents the distance between two nodes. The Spatial Temporal Reach and Escape Logic (STREL) [13], [14] generalizes SSSL, by introducing two new spatial operators, (*reach* and *escape*), which are able to express the same spatial operators of SSSL. Furthermore, while SSSL can be applied only on static weight undirected graphs, STREL can be applied also to dynamic networks. However, both SSSL and STREL do not support spatial aggregation operators that we show to be an important feature for monitoring smart cities.

## X. CONCLUSION

In this paper, we present a novel Spatial Aggregation Signal Temporal Logic to specify and to monitor requirements of smart cities at runtime. We develop an efficient monitoring

framework that optimizes the requirement parsing process and can check in parallel a SaSTL requirement over multiple data streams generated from thousands of sensors that are typically spatially distributed over a smart city. SaSTL is a powerful specification language for smart cities because of its capability to monitor the city desirable features of temporal (e.g., interval), spatial (e.g., PoIs, range) and their complicated relations (e.g. always, everywhere, aggregation) between them. More importantly, it can coalesce many requirements into a single SaSTL formula and provide the aggregated results efficiently, which is a major advance on what smart cities do now. The development of 5G and 6G could better support the monitoring and communication among sensors, services and the city center. We believe it is a valuable step towards developing a practical smart city monitoring system even though there are still open issues for future work. Furthermore, SaSTL monitor can also be easily generalized and applied to monitor other large-scale IoT deployments at runtime efficiently. In the future, we will explore its capability to specify and monitor other properties and requirements (e.g., security and privacy).

## ACKNOWLEDGEMENT

This work was supported in part by National Science Foundation grants CCF-1942836, CNS-1952096 and by the Austrian FFG-funded IoT4CPS project at TU Wien.

## REFERENCES

- [1] C. E. Catlett, P. H. Beckman, R. Sankaran, and K. K. Galvin, "Array of things: a scientific research instrument in the public way: platform design and early lessons learned," in *Proceedings of the 2nd International Workshop on Science of Smart City Operations and Platforms Engineering*. ACM, 2017, pp. 26–33.
- [2] New York Times, "IBM takes 'smarter cities' to rio de janeiro," 2012.
- [3] Cisco, "Smart+connected operations center," 2017.
- [4] M. Ma, S. M. Preum, M. Y. Ahmed, W. Tärneberg, A. Hendawi, and J. A. Stankovic, "Data sets, modeling, and decision making in smart cities: A survey," *ACM Transactions on Cyber-Physical Systems*, vol. 4, no. 2, pp. 1–28, 2019.
- [5] M. Ma, S. M. Preum, and J. A. Stankovic, "Cityguard: A watchdog for safety-aware conflict detection in smart cities," in *Proceedings of the Second International Conference on Internet-of-Things Design and Implementation*, 2017, pp. 259–270.
- [6] H. Zhang, Y. Zheng, and Y. Yu, "Detecting urban anomalies using multiple spatio-temporal data sources," *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 2, no. 1, p. 54, 2018.
- [7] S. Sheng, E. Pakdamanian, K. Han, B. Kim, P. Tiwari, I. Kim, and L. Feng, "A case study of trust on autonomous driving," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 4368–4373.
- [8] M. Ma, J. A. Stankovic, and L. Feng, "Runtime monitoring of safety and performance requirements in smart cities," in *1st ACM Workshop on the Internet of Safe Things*, 2017.

- [9] I. Haghghi, A. Jones, Z. Kong, E. Bartocci, R. Gros, and C. Belta, "Spatel: a novel spatial-temporal logic and its applications to networked systems," in *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*. ACM, 2015, pp. 189–198.
- [10] M. Ma, J. A. Stankovic, and L. Feng, "Cityresolver: a decision support system for conflict resolution in smart cities," in *Proceedings of the 9th ACM/IEEE International Conference on Cyber-Physical Systems*. IEEE Press, 2018, pp. 55–64.
- [11] O. Maler and D. Nickovic, "Monitoring temporal properties of continuous signals," in *Proc. FORMATS*, 2004.
- [12] L. Nenzi, L. Bortolussi, V. Ciancia, M. Loreti, and M. Massink, "Qualitative and quantitative monitoring of spatio-temporal properties," in *Runtime Verification - 6th International Conference, RV 2015*, vol. 9333. Springer, 2015, pp. 21–37.
- [13] E. Bartocci, L. Bortolussi, M. Loreti, and L. Nenzi, "Monitoring mobile and spatially distributed cyber-physical systems," in *MEMOCODE 2017: the 15th ACM-IEEE International Conference on Formal Methods and Models for System Design*. ACM, 2017, pp. 146–155.
- [14] E. Bartocci, L. Bortolussi, M. Loreti, L. Nenzi, and S. Silvetti, "Moonlight: A lightweight tool for monitoring spatio-temporal properties," in *Proc. of RV 2020: the 20th International Conference on Runtime Verification*, ser. LNCS, vol. 12399. Springer, 2020, pp. 417–428.
- [15] L. Nenzi, E. Bartocci, L. Bortolussi, M. Loreti, and E. Visconti, "Monitoring spatio-temporal properties (invited tutorial)," in *Proc. of RV 2020: the 20th International Conference on Runtime Verification*, ser. LNCS, vol. 12399. Springer, 2020, pp. 21–46.
- [16] M. Ma, E. Bartocci, E. Lifland, J. Stankovic, and L. Feng, "Sastl: Spatial aggregation signal temporal logic for runtime monitoring in smart cities," in *2020 ACM/IEEE 11th International Conference on Cyber-Physical Systems (ICCPs)*. IEEE, 2020, pp. 51–62.
- [17] NYC.gov, "Emissions from transportation, nyc environment protection," 2019. [Online]. Available: [https://www1.nyc.gov/html/dep/html/air/emissions\\_from\\_transportation.shtml](https://www1.nyc.gov/html/dep/html/air/emissions_from_transportation.shtml)
- [18] District of Columbia Municipal Regulations and D. of Columbia Register, "Air quality - motor vehicular pollutants, lead, odors, and nuisance pollutants," 2016.
- [19] S. Matteo and J. Brannan, "A local law to amend the administrative code of the city of new york, in relation to restricting the use of bus lanes by sight-seeing buses," in *Restricting the use of bus lanes by sight-seeing buses*. The New York City Council, 2019.
- [20] NYC Environment Protection, "Use of heating oil remaining in tanks." The city of New York, 2019.
- [21] United States Environmental Protection Agency, "Residential energy efficiency," in *Energy Resources for State and Local Governments*. The city of New York, 2019.
- [22] LA Sec 111.03. Minimum Ambient Noise Level, "Official city of los angeles municipal code," 2016.
- [23] Hong Kong, "Guide to indoor air quality management in hong kong regional offices and public places," in *Guide to Indoor Air Quality Management*, 2019.
- [24] NYC.gov, "Stopping, standing or parking prohibited in specified places," in *New York Public Law*, 2016.
- [25] Beijing Emergency Agency, "Pre-hospital medical emergency regulations," 2016.
- [26] Beijing Government, "Safety management for kindergarten, primary and secondary school," 2016.
- [27] R. E. Ladner and M. J. Fischer, "Parallel prefix computation," *Journal of the ACM (JACM)*, vol. 27, no. 4, pp. 831–838, 1980.
- [28] A. Donz , T. Ferrere, and O. Maler, "Efficient robust monitoring for STL," in *International Conference on Computer Aided Verification*. Springer, 2013, pp. 264–279.
- [29] G. S. Lueker, "A data structure for orthogonal range queries," in *19th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, 1978, pp. 28–34.
- [30] City of Chicago, "Crimes of Chicago - one year prior to present," <https://data.cityofchicago.org/Public-Safety/Crimes-Map/dfnk-7re6>, 2018.
- [31] M. Ma, S. M. Preum, W. Tarneberg, M. Ahmed, M. Ruiters, and J. Stankovic, "Detection of runtime conflicts among services in smart cities," in *2016 IEEE International Conference on Smart Computing (SMARTCOMP)*. IEEE, 2016, pp. 1–10.
- [32] M. Behrisch, L. Bieker, J. Erdmann, and D. Krajzewicz, "Sumo-simulation of urban mobility: an overview," in *Proceedings of SIMUL 2011*. ThinkMind, 2011.
- [33] NYC.gov, *New York City Open Data*, <https://nycopendata.socrata.com/>.
- [34] S. Bischof, A. Karapantelakis, C.-S. Nechifor, A. P. Sheth, A. Mileo, and P. Barnaghi, "Semantic modelling of smart city data," 2014.
- [35] C. L. Talcott, "Cyber-physical systems and events," in *Software-Intensive Systems and New Computing Paradigms - Challenges and Visions*, ser. LNCS. Springer, 2008, vol. 5380, pp. 101–115.
- [36] Y. Tan, M. C. Vuran, and S. Goddard, "Spatio-temporal event model for cyber-physical systems," in *2009 29th IEEE International Conference on Distributed Computing Systems Workshops*. IEEE, 2009, pp. 44–50.
- [37] L. Caires and L. Cardelli, "A spatial logic for concurrency (part i)," *Information and Computation*, vol. 186, no. 2, pp. 194 – 235, 2003.
- [38] F. B. Buonamici, G. Belmonte, V. Ciancia, D. Latella, and M. Massink, "Spatial logics and model checking for medical imaging," *Int. J. Softw. Tools Technol. Transf.*, vol. 22, no. 2, pp. 195–217, 2020.
- [39] B. Bennett, A. G. Cohn, F. Wolter, and M. Zakharyashev, "Multi-dimensional modal logic as a framework for spatio-temporal reasoning," *Applied Intelligence*, vol. 17, no. 3, pp. 239–251, Sep. 2002.
- [40] D. Bresolin, P. Sala, D. D. Monica, A. Montanari, and G. Sciavicco, "A decidable spatial generalization of metric interval temporal logic," in *2010 17th International Symposium on Temporal Representation and Reasoning*, 2010, pp. 95–102.
- [41] M. Marx and M. Reynolds, "Undecidability of compass logic," *J Logic Computation*, vol. 9, no. 6, pp. 897–914, 1999.
- [42] I. Haghghi, A. Jones, J. Z. Kong, E. Bartocci, G. R., and C. Belta, "SpaTeL: A Novel Spatial-Temporal Logic and Its Applications to Networked Systems," in *Proc. of HSCC*, 2015.
- [43] E. A. Gol, E. Bartocci, and C. Belta, "A formal methods approach to pattern synthesis in reaction diffusion systems," in *Proc. of CDC*, 2014.
- [44] E. Bartocci, E. A. Gol, I. Haghghi, and C. Belta, "A formal methods approach to pattern recognition and synthesis in reaction diffusion networks," *IEEE Trans. Control. Netw. Syst.*, vol. 5, no. 1, pp. 308–320, 2018.
- [45] R. Grosu, S. A. Smolka, F. Corradini, A. Wasilewska, E. Entcheva, and E. Bartocci, "Learning and detecting emergent behavior in networks of cardiac myocytes," *Commun. ACM*, vol. 52, no. 3, pp. 97–105, 2009. [Online]. Available: <http://doi.acm.org/10.1145/1467247.1467271>
- [46] E. Bartocci, F. Corradini, M. R. D. Berardini, E. Entcheva, S. A. Smolka, and R. Grosu, "Modeling and simulation of cardiac tissue using hybrid I/O automata," *Theor. Comput. Sci.*, vol. 410, no. 33–34, pp. 3149–3165, 2009.

**Meiyi Ma** is a Ph.D. candidate of Computer Science at the University of Virginia. Her research interests are at the intersection of cyber-physical systems, deep learning and formal methods.

**Ezio Bartocci** is a full professor at the Faculty of Computer Science of TU Wien, where he leads the Trustworthy Cyber-Physical Systems (TrustCPS) Group. The primary focus of his research is to develop formal methods, computational tools and techniques that support the modeling and the automated analysis of complex computational systems, including software systems, cyber-physical systems and biological systems.

**Eli Lifland** is a software engineer at Ought. He received his Bachelor's degree in Computer Science and Economics at the University of Virginia.

**John A. Stankovic** is the BP America Professor in the Computer Science Department at the University of Virginia and Director of the Link Lab. He is a Fellow of both the IEEE and the ACM. He has been awarded an Honorary Doctorate from the University of York for his work on real-time systems. His research interests are in smart and connected health, cyber physical systems, and the Internet of Things. Prof. Stankovic received his PhD from Brown University.

**Lu Feng** is an Assistant Professor of Computer Science at the University of Virginia. Her research interests are in cyber-physical systems and formal methods. Dr. Feng received her PhD in Computer Science from the University of Oxford in 2014. She is a member of ACM and IEEE.

## APPENDIX

## 1. Preliminaries on Signal Temporal Logic

The syntax of an STL formula  $\varphi$  is usually defined as follows,

$$\varphi ::= \mu \mid \neg\varphi \mid \varphi \wedge \varphi \mid \diamond_{(a,b)}\varphi \mid \square_{(a,b)}\varphi \mid \varphi \mathbf{U}_{(a,b)}\varphi.$$

We call  $\mu$  a signal predicate, which is a formula in the form of  $f(x) \geq 0$  with a signal variable  $x \in \mathcal{X}$  and a function  $f: \mathcal{X} \rightarrow \mathbb{R}$ . The temporal operators  $\square$ ,  $\diamond$ , and  $\mathbf{U}$  denote ‘‘always’’, ‘‘eventually’’ and ‘‘until’’, respectively. The bounded interval  $(a, b)$  denotes the time interval of temporal operators.

Below we present the formal definition of STL Boolean semantics. To informally explain the STL operations, formula  $\square_{(a,b)}\varphi$  is true iff  $\varphi$  is always true in the time interval  $(a, b)$ . Formula  $\diamond_{(a,b)}\varphi$  is true iff  $\varphi$  is true at sometime between  $a$  and  $b$ . Formula  $\varphi_1 \mathbf{U}_{(a,b)}\varphi_2$  is true iff  $\varphi_1$  is true until  $\varphi_2$  becomes true at sometime between  $a$  and  $b$ .

$$\begin{aligned} (\omega, t) \models \mu & \Leftrightarrow f(x) > 0 \\ (\omega, t) \models \neg\varphi & \Leftrightarrow (\omega, t) \not\models \varphi \\ (\omega, t) \models \varphi_1 \wedge \varphi_2 & \Leftrightarrow (\omega, t) \models \varphi_1 \text{ and } (\omega, t) \models \varphi_2 \\ (\omega, t) \models \square_{(a,b)}\varphi & \Leftrightarrow \forall t \in (a, b), (\omega, t) \models \varphi \\ (\omega, t) \models \diamond_{(a,b)}\varphi & \Leftrightarrow \exists t \in (a, b) \cap \mathbb{T}, (\omega, t) \models \varphi \\ (\omega, t) \models \varphi_1 \mathbf{U}_{(a,b)}\varphi_2 & \Leftrightarrow \exists t' \in (t+a, t+b) \cap \mathbb{T}, (\omega, t') \models \varphi_2 \\ & \text{and } \forall t'' \in (t, t'), (\omega, t'') \models \varphi_1 \end{aligned}$$

Next, we present the formal definition of STL quantitative semantics.

$$\begin{aligned} \rho(x \sim c, \omega, t) &= \pi_x(\omega)[t] - c \\ \rho(\neg\varphi, \omega, t) &= -\rho(\varphi, \omega, t) \\ \rho(\varphi_1 \wedge \varphi_2, \omega, t) &= \min\{\rho(\varphi_1, \omega, t), \rho(\varphi_2, \omega, t)\} \\ \rho(\square_I\varphi, \omega, t) &= \min_{t' \in (t, t+I)} \rho(\varphi, \omega, t') \\ \rho(\diamond_I\varphi, \omega, t) &= \max_{t' \in (t, t+I)} \rho(\varphi, \omega, t') \\ \rho(\varphi_1 \mathbf{U}_I \varphi_2, \omega, t) &= \sup_{t' \in (t+I) \cap \mathbb{T}} (\min\{\rho(\varphi_2, \omega, t'), \\ & \inf_{t'' \in [t, t']} (\rho(\varphi_1, \omega, t''))\}) \end{aligned}$$

## 2. Proofs

**Theorem 1** (Soundness, restate). *Let  $\varphi$  be an STL formula,  $\omega$  a trace and  $t$  a time,*

$$\begin{aligned} \rho(\varphi, \omega, t, l) > 0 & \Rightarrow (\omega, t, l) \models \varphi \\ \rho(\varphi, \omega, t, l) < 0 & \Rightarrow (\omega, t, l) \not\models \varphi \end{aligned}$$

*Proof.* We prove the first property  $\rho(\varphi, \omega, t, l) > 0 \Rightarrow (\omega, t, l) \models \varphi$  by induction:

First we show the soundness property hold for the predicate  $\varphi := \mu$ . In this case, we have  $\rho(\varphi, \omega, t, l) = f(x)$ . Therefore, if  $\rho(\varphi, \omega, t, l) > 0$  we have  $f(x) > 0$ , that is,  $(\omega, t, l) \models \varphi$ .

Case  $\varphi := \neg\varphi'$ : We have  $\rho(\varphi, \omega, t) = -\rho(\varphi', \omega, t, l) > 0$ . Therefore we have  $\rho(\varphi', \omega, t, l) < 0$ , that is,  $(\omega, t, l) \not\models \varphi'$ , which is equivalent to  $(\omega, t, l) \models \varphi$  by definition.

Case  $\varphi := \varphi_1 \wedge \varphi_2$ : We have  $\rho(\varphi_1 \wedge \varphi_2, \omega, t, l) = \min\{\rho(\varphi_1, \omega, t, l), \rho(\varphi_2, \omega, t, l)\} > 0$ . Therefore, we have

$\rho(\varphi_1, \omega, t, l) > 0$  and  $\rho(\varphi_2, \omega, t, l) > 0$ . Thus,  $(\omega, t, l) \models \varphi_1$  and  $(\omega, t, l) \models \varphi_2$ . By definition, we have  $(\omega, t, l) \models \varphi$ .

Case  $\varphi := \varphi_1 \mathbf{U}_I \varphi_2$ :  $\rho(\varphi_1 \mathbf{U}_I \varphi_2, \omega, t, l) = \max_{t' \in (t, t+I)} \{\min\{\rho(\varphi_2, \omega, t', l), \min_{t'' \in (t, t')} \rho(\varphi_1, \omega, t'', l)\}\} > 0$ . We have  $\exists t' \in (t+I), \min\{\rho(\varphi_2, \omega, t', l), \min_{t'' \in (t, t')} \rho(\varphi_1, \omega, t'', l)\} > 0$ . Therefore,  $\exists t' \in (t+I), \rho(\varphi_2, \omega, t', l) > 0 \wedge \min_{t'' \in (t, t')} \rho(\varphi_1, \omega, t'', l) > 0$ . Thus, it's equivalent to  $\exists t' \in (t+I) \cap \mathbb{T}, (\omega, t', l) \models \varphi_2$  and  $\forall t'' \in (t, t'), (\omega, t'', l) \models \varphi_1$ . By definition, we have  $(\omega, t, l) \models \varphi$ .

Case  $\varphi := \mathcal{A}_D^{\text{op}} x \sim c$ : we have  $\rho(\mathcal{A}_D^{\text{op}} x \sim c, \omega, t, l) > 0$ , which indicates  $\text{op}(\alpha_D^x(\omega, t, l)) - c > 0$ , following the definition, we have  $(\mathcal{A}_D^{\text{op}} x \sim c, \omega, t, l) \models \varphi$ .

Case  $\varphi := \mathcal{C}_D^{\text{op}} \varphi \sim c$  when  $\text{op} = \max$ , we have  $\max_{l' \in L_D^l} \{\rho(\varphi, \omega, t, l')\} > 0$ , thus, there is at least one location  $l \in \mathcal{D}, \rho(\varphi, \omega, t, l) > 0$ , i.e.,  $(\omega, t, l) \models \varphi$ , therefore,  $\max(\{g((\omega, t, l') \models \varphi) \mid l' \in L_D^l\}) > c$  ( $c \in [0, 1)$ ) is true, therefore,  $(\omega, t, l) \models \mathcal{C}_D^{\max} \varphi > c$ . when  $\text{op} = \min$ , we have  $\min_{l' \in L_D^l} \{\rho(\varphi, \omega, t, l')\} > 0$ , thus, for any location,  $\rho(\varphi, \omega, t, l) > 0$ , i.e.,  $l \in \mathcal{D}, (\omega, t, l) \models \varphi$ , therefore,  $\min(\{g((\omega, t, l') \models \varphi) \mid l' \in L_D^l\}) > c$  ( $c \in [0, 1)$ ) is true, therefore,  $(\omega, t, l) \models \mathcal{C}_D^{\min} \varphi \sim c$ . When  $\text{op} = \text{sum}$ , we have  $\delta([c], \{\rho(\varphi, \omega, t, l') \mid l' \in L_D^l\}) > 0$ , thus, for at least  $[c]$  locations  $l$ , we have  $\rho(\varphi, \omega, t, l) \mid l \in L_D^l > 0$ , i.e.,  $\text{sum}(\{g((\omega, t, l) \models \varphi) \mid l \in L_D^l\}) > c$  is true, therefore,  $(\omega, t, l) \models \mathcal{C}_D^{\text{sum}} \varphi > c$ . Similarly, we can prove when  $\text{op} = \text{avg}$ , if  $\delta([c \times |L_D^l|], \{\rho(\varphi, \omega, t, l') \mid l' \in L_D^l\}) > 0$ , then  $(\omega, t, l) \models \mathcal{C}_D^{\text{avg}} \varphi > c$ .  $\square$

**Theorem 2** (Correctness, restate). *Let  $\varphi$  be an STL formula,  $\omega$  and  $\omega'$  traces over the same time and spatial domains, and  $t, l \in \text{dom}(\varphi, \omega)$ , then*

$$(\omega, t, l) \models \varphi \text{ and } \|\omega - \omega'\|_\infty < \rho(\varphi, \omega, t, l) \Rightarrow (\omega', t, l) \models \varphi$$

*Proof.* First, whenever  $\rho(\varphi, \omega, t, l) \neq 0$ , its sign indicates the satisfaction status.

By induction, we have the following cases:

Case  $\varphi := x \sim c$ : We have  $\rho(\varphi, \omega', t, l) = \pi_x(\omega')[t, l] - c \geq \pi_x(\omega)[t, l] - c - \|\omega - \omega'\|_\infty = \rho(\varphi, \omega, t, l) - \|\omega - \omega'\|_\infty > 0$ . Therefore, we have  $(\omega', t, l) \models \varphi$ .

Case  $\varphi := \neg\varphi'$ : We have  $\rho(\varphi, \omega', t, l) = -\rho(\varphi', \omega', t, l)$ . By the inductive assumption we have  $\rho(\varphi', \omega', t, l) < 0$ . Therefore, we have  $(\omega', t, l) \models \varphi$ .

Case  $\varphi := \varphi_1 \vee \varphi_2$ : Following the condition, we have either  $(\omega, t, l) \models \varphi_1$  holds or  $(\omega, t, l) \models \varphi_2$  holds. We also have  $\rho(\varphi, \omega', t, l) = \max\{\rho(\varphi_1, \omega', t, l), \rho(\varphi_2, \omega', t, l)\}$ . If  $(\omega, t, l) \models \varphi_1$ , by the inductive assumption we have  $\rho(\varphi_1, \omega', t, l) > 0$ . Therefore,  $\rho(\varphi, \omega', t, l) > 0$ . Similarly, if  $(\omega, t, l) \models \varphi_2$ , by the inductive assumption we have  $\rho(\varphi_2, \omega', t, l) > 0$ . Therefore, we have  $(\omega', t, l) \models \varphi$ .

Case  $\varphi := \varphi_1 \mathbf{U}_I \varphi_2$ : As  $(\omega, t, l) \models \varphi$ , there exists  $t'$  that  $\forall t'' \in (t, t'), \rho(\varphi_1, \omega, t'', l) \geq \rho(\varphi, \omega, t, l)$  and  $\rho(\varphi_2, \omega, t', l) \geq \rho(\varphi, \omega, t, l)$ . By the inductive assumption, we have  $(\omega', t', l) \models \varphi_2$  and  $\forall t'' \in (t, t'), (\omega', t'', l) \models \varphi_1$ . Therefore, we have  $(\omega', t, l) \models \varphi$ .

Case  $\varphi := \mathcal{A}_D^{\text{op}} x \sim c$ :

- When  $op = \text{sum}$ ,  $\rho(\phi, \omega', t, l) = \frac{\text{sum}(\alpha_{\mathcal{D}}^x(\omega', t, l)) - c}{|\alpha_{\mathcal{D}}^x(\omega', t, l)|} \geq \frac{\text{sum}(\alpha_{\mathcal{D}}^x(\omega, t, l)) - c - \sum_{d \in \alpha_{\mathcal{D}}^x(\omega, t, l)} \|\omega - \omega'\|_{\infty}}{|\alpha_{\mathcal{D}}^x(\omega, t, l)|} = \rho(\phi, \omega, t, l) - \|\omega - \omega'\|_{\infty} > 0$ . Therefore, we have  $(\omega', t, l) \models \phi$ .
- When  $op \neq \text{sum}$ , we first show that  $op(\alpha_{\mathcal{D}}^x(\omega, t, l)) - op(\alpha_{\mathcal{D}}^x(\omega', t, l)) \leq \|\omega - \omega'\|_{\infty}$ . Recall the definition that  $\alpha_{\mathcal{D}}^x(\omega, t, l) := \{\pi_x(\omega)[t, l'] \mid l' \in L_{\mathcal{D}}^l \text{ and } \pi_x(\omega)[t, l'] \neq \perp\}$ . For any combination of  $t$  and  $l$ ,  $\pi_x(\omega)[t, l] \leq \pi_x(\omega')[t, l] + \|\omega - \omega'\|_{\infty}$ . As all the items of  $\alpha_{\mathcal{D}}^x(\omega, t, l)$  holds the property, for the operators  $\text{max}$ ,  $\text{min}$  and  $\text{avg}$ ,  $op(\alpha_{\mathcal{D}}^x(\omega, t, l)) - op(\alpha_{\mathcal{D}}^x(\omega', t, l)) \leq \|\omega - \omega'\|_{\infty}$ . Therefore we have  $\rho(\phi, \omega', t, l) = op(\alpha_{\mathcal{D}}^x(\omega', t, l)) - c \geq op(\alpha_{\mathcal{D}}^x(\omega, t, l)) - \|\omega - \omega'\|_{\infty} - c = \rho(\phi, \omega, t, l) - \|\omega - \omega'\|_{\infty} > 0$ , which indicates  $(\omega', t, l) \models \phi$ .

Case  $\varphi = C_{\mathcal{D}}^{\text{op}} \varphi' \sim c$ :

- When  $op = \text{sum}$ , as  $\rho(C_{\mathcal{D}}^{\text{op}} \varphi' \sim c, \omega, t, l) = \delta([\![c]\!] , \{\rho(\varphi', \omega', t, l') \mid l' \in L_{\mathcal{D}}^l\})$ , we know that there exists at least  $\lceil c \rceil$  different  $l' \in L_{\mathcal{D}}^l$  that  $\rho(\varphi', \omega, t, l') \geq \rho(C_{\mathcal{D}}^{\text{op}} \varphi' \sim c, \omega, t, l) > \|\omega - \omega'\|_{\infty}$ . By the inductive rule, we have at least  $\lceil c \rceil$  different  $l' \in L_{\mathcal{D}}^l$  that  $\rho(\varphi', \omega', t, l') > 0$ . Therefore, by the definition of  $\rho(C_{\mathcal{D}}^{\text{op}} \varphi' \sim c)$  of we have  $(\omega', t, l) \models \phi$ .
- Similarly when  $op = \text{avg}$ , as  $\rho(C_{\mathcal{D}}^{\text{op}} \varphi' \sim c, \omega, t, l) = \delta([\![c \times |L_{\mathcal{D}}^l|\!] \rceil \{\rho(\varphi', \omega', t, l') \mid l' \in L_{\mathcal{D}}^l\})$ , we know that there exists at least  $\lceil c \times |L_{\mathcal{D}}^l|\rceil$  different  $l' \in L_{\mathcal{D}}^l$  that  $\rho(\varphi', \omega, t, l') \geq \rho(C_{\mathcal{D}}^{\text{op}} \varphi' \sim c, \omega, t, l) > \|\omega - \omega'\|_{\infty}$ . By the inductive rule, we have at least  $\lceil c \times |L_{\mathcal{D}}^l|\rceil$  different  $l' \in L_{\mathcal{D}}^l$  that  $\rho(\varphi', \omega', t, l') > 0$ . Therefore, we have  $(\omega', t, l) \models \phi$ .
- When  $op = \text{max}$ ,  $\rho(C_{\mathcal{D}}^{\text{op}} \varphi' \sim c, \omega, t, l) = \max_{l' \in L_{\mathcal{D}}^l} \{\rho(\varphi, \omega, t, l')\}$ . Let  $l'$  be the location that  $\rho(\varphi, \omega, t, l')$  achieves maximum, we have  $\rho(\varphi', \omega, t, l') \geq \rho(C_{\mathcal{D}}^{\text{op}} \varphi' \sim c, \omega, t, l) > \|\omega - \omega'\|_{\infty}$ . By the inductive rule,  $\rho(\varphi', \omega', t, l') > 0$ . Therefore, we have  $(\omega', t, l) \models \phi$ .
- When  $op = \text{min}$ ,  $\rho(C_{\mathcal{D}}^{\text{op}} \varphi' \sim c, \omega, t, l) = \min_{l' \in L_{\mathcal{D}}^l} \{\rho(\varphi, \omega, t, l')\}$ . We have for every  $l' \in L_{\mathcal{D}}^l$ ,  $\rho(\varphi', \omega, t, l') \geq \rho(C_{\mathcal{D}}^{\text{op}} \varphi' \sim c, \omega, t, l) > \|\omega - \omega'\|_{\infty}$ . By the inductive rule, We have for every  $l' \in L_{\mathcal{D}}^l$  that  $\rho(\varphi', \omega', t, l') > 0$ . Therefore, we have  $(\omega', t, l) \models \phi$ .  $\square$

**Lemma 3** (Complexity of spatial operators, restate). *The time complexity to monitor at each location  $l$  at time  $t$  the satisfaction of a spatial operator such as  $\boxplus_{\mathcal{D}}$ ,  $\boxtimes_{\mathcal{D}}$ ,  $\mathcal{A}_{\mathcal{D}}^{\text{op}}$ , and  $C_{\mathcal{D}}^{\text{op}}$  is  $O(\log(n) + |L|)$  where  $L$  is the set of locations at distance within the range  $\mathcal{D}$  from  $l$ .*

*Proof.* According to [29], the time complexity to retrieve a set of nodes  $L$  with a distance to a desired location in a range  $\mathcal{D}$  from a location  $l$  is  $O(\log(n) + |L|)$ . The aggregation and counting operations of Algorithm 7 and Algorithm 8 can be performed while the locations are retrieved.  $\square$

**Theorem 6.** *The time complexity of the SaSTL monitoring algorithm is upper-bounded by  $O(|\phi| \times T_{\text{max}} \times (\log(n) + |L|_{\text{max}}))$  where  $T_{\text{max}}$  is the largest number of samples of the intervals considered in the temporal operators of  $\phi$  and  $|L|_{\text{max}}$  is the*

---

**Algorithm 6:** SaSTL Boolean monitoring algorithm  
MonitorB( $\varphi, \omega, t, l, G$ )

---

**Input :** SaSTL Requirement  $\varphi$ , Signal  $\omega$ , Time  $t$ , Location  $l$ , weighted undirected graph  $G$

**Output:** Boolean Satisfaction Value

```

begin
  switch  $\varphi$  do
    Case  $x \sim c$ 
      | return  $\pi_x(\omega)[t, l] \sim c$ ;
    Case  $\neg\varphi$ 
      | return  $\neg$  MonitorB( $\varphi, \omega, t, l, G$ );
    Case  $\varphi_1 \wedge \varphi_2$ ;
      | See Alg. 4 for an update
      | return MonitorB( $\varphi_1, \omega, t, l, G$ )  $\wedge$  MonitorB( $\varphi_2, \omega, t, l, G$ )
    Case  $\varphi_1 \cup_I \varphi_2$ 
      | Boolean  $f := \text{True}$ ;
      | for  $t' \in (t + I) \cap \mathbb{T}$  do
          | if Monitor( $\varphi_2, \omega, t', l, G$ ) then
              |  $f := \text{True}$ ;
              | for  $t'' \in [t, t']$  do
                  |  $f := f \wedge$  Monitor( $\varphi_1, \omega, t'', l, G$ );
                  | if ( $\neg f$ ) then break;
              | end
          | end
          | if ( $f$ ) then return True;
      | end
      | return False;
    Case  $\mathcal{A}_{\mathcal{D}}^{\text{op}} x \sim c$ ;
      | See Alg. 7
      | return AggregateB( $x, c, op, \mathcal{D}, t, l, G$ );
    Case  $C_{\mathcal{D}}^{\text{op}} \varphi \sim c$ ;
      | See Alg. 8 and Alg. 5
      | return CountingNeighboursB( $\varphi, c, op, \mathcal{D}, t, l, G$ );
  end
end

```

---



---

**Algorithm 7:** AggregateB( $x, op, \mathcal{D}, \omega, t, l, G$ )

---

**Function** AggregateB( $x, c, op, \mathcal{D}, \omega, t, l, G$ ):

```

begin
  Real  $v := 0$ ;  $n := 0$ ;
  if  $op == \text{"min"}$  then  $v := \infty$ ;
  if  $op == \text{"max"}$  then  $v := -\infty$ ;
   $L_{\mathcal{D}}^l := \text{deScan}(l, G, \mathcal{D})$ 
  for  $l' \in L_{\mathcal{D}}^l$  do
    | if  $op \in \{\text{min}, \text{max}, \text{sum}\}$  then
        |  $v := op(v, \pi_x(\omega)[t, l'])$ ;
    | end
    | if  $op == \text{"avg"}$  then
        |  $v := \text{sum}(v, \pi_x(\omega)[t, l'])$ ;
    | end
    |  $n := n + 1$ 
  end
  if  $op == \text{"avg"} \wedge n \neq 0$  then  $v := v/n$ ;
  if  $n == 0$  then
    | return True
  else
    | return  $v \sim c$ ;
  end
end

```

---

maximum number of locations defined by the spatial temporal operators of  $\phi$ .

*Proof.* Following Lemma 3, by considering  $T_{\text{max}}$  the worst possible number of samples that we need to consider for all possible intervals of temporal operators present in the formula, and  $|L|_{\text{max}}$  for the worst possible number of locations that we need to consider for all possible intervals of spatial operators present in the formula. When there are two or more operators nested, the time complexity for one operation is bounded by  $O(T_{\text{max}} (\log(n) + |L|_{\text{max}}))$ . As there are  $|\phi|$  nodes in the syntax tree of  $\phi$ , the time complexity of the SaSTL monitoring algorithm is bounded by the summation over all  $|\phi|$  nodes, which is  $O(|\phi| T_{\text{max}} (\log(n) + |L|_{\text{max}}))$ .  $\square$



Table VII: List of services running in simulated NYC

Service	Description
<b>S1:</b> Traffic Service	It controls traffic signals in street intersections to relieve congestion and optimize or improve traffic performance.
<b>S2:</b> Emergency Service	It requests green traffic signals in order to transport patients in critical conditions to hospitals as soon as possible.
<b>S3:</b> Accident Service	It blocks a street where some accident occurs and alert nearby vehicles to detour.
<b>S4:</b> Infrastructure Service	It schedules infrastructure check-up and repair appointments.
<b>S5:</b> Pedestrian Service	It shortens the pedestrians' waiting time by adjusting traffic signals when pedestrians wait in the intersection.
<b>S6:</b> Air Pollution Control	It adjusts the traffic by adjusting traffic signal and sending speed request to vehicles when CO emission is high.
<b>S7:</b> PM2.5/ PM10 Control	It adjusts the traffic when PM2.5/ PM10 emission is high by adjusting traffic signal and sending speed request to vehicles directly.
<b>S8:</b> Parking Service	It directs the driver to the nearest parking lot.
<b>S9:</b> Noise Control	When noise level exceeds its threshold, it controls the number of vehicles going through related streets and redirect vehicles on the streets by adjusting traffic signals.
<b>S10:</b> Event Service	It ensures operation of a city event by blocking the lanes nearby the event.

**Algorithm 8:** CountingNeighboursB( $x, op, \mathcal{D}, \omega, t, l, G$ )

```

begin
  Real  $v := 0; n := 0$ 
  if  $op == "min"$  then  $v := \infty$ ;
  if  $op == "max"$  then  $v := -\infty$ ;
   $L_{\mathcal{D}}^l := \text{deScan}(l, G, \mathcal{D})$ 
  for  $l' \in L_{\mathcal{D}}^l$  do
    if  $\text{Monitor}(\varphi, \omega, t, l, G) \wedge op \in \{min, max, sum\}$  then
      |  $v := op(v, 1)$ ;
    end
    if  $\text{Monitor}(\varphi, \omega, t, l, G) \wedge op == "avg"$  then
      |  $v := sum(v, 1)$ ;
    end
     $n := n + 1$ 
  end
  if  $op == "avg" \wedge n \neq 0$  then  $v := v/n$ ;
  if  $n == 0$  then
    | return True
  else
    | return  $v \sim c$ ;
  end
end

```

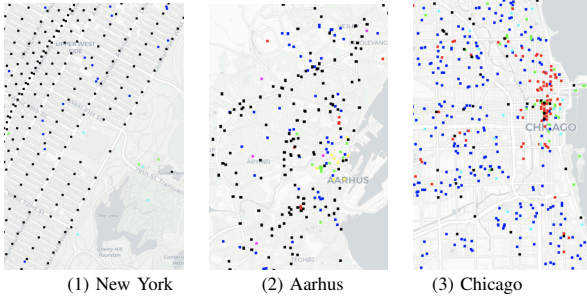


Figure 10: Partial Maps of Chicago, Aarhus and New York with POIs and sensors annotated. (The black nodes represent the locations of sensors, red nodes represent the locations of hospitals, dark blue nodes represent schools, light blue nodes represent parks and green nodes represent theaters.)

**3. Monitoring Algorithms** We presented the details of the Boolean monitoring algorithms in Algorithm 6 with Algorithm 7 for the aggregation operation and Algorithm 8 for the counting operation.

**4. Smart Services in Simulated NYC** In the evaluation section, we set up the simulator with ten smart services. The description of these services are presented in Table VII.