

Cloud Deployment Tradeoffs for the Analysis of Spatially-Distributed Systems of Internet-of-Things

CHRISTOS TSIGKANOS, Technische Universität Wien, Austria

MARTIN GARRIGA, Universidad Nacional del Comahue, Argentina

LUCIANO BARESI, Politecnico di Milano, Italy

CARLO GHEZZI, Politecnico di Milano, Italy

Internet-enabled things and devices operating in the physical world are increasingly integrated in modern distributed systems, supporting functionalities that require assurances that certain critical requirements are satisfied by the overall system. We focus here on spatially-distributed Internet-of-Things systems such as smart environments, where the dynamics of spatial distribution of entities in the system is crucial to requirements satisfaction. Analysis techniques need to be in place while systems operate to ensure that requirements are fulfilled. This may be achieved by keeping a model of the system at runtime, monitoring events that lead to changes in the spatial environment, and performing analysis. This computationally-intensive runtime assurance method cannot be supported by resource-constrained devices that populate the space and must be offloaded to the cloud. However, challenges arise regarding resource allocation and cost, especially when the workload is unknown at the system's design time. As such, it may be difficult or even impossible to guarantee application service level agreements, e.g., on response times. To this end, we instantiate spatial verification processes, integrating them to the service layer of an IoT-cloud architecture based on microservices. We propose several cloud deployments for such an architecture for assurance of spatial requirements – based on virtual machines, containers, and the recent Functions-as-a-Service paradigm. Then, we assess deployments' tradeoffs in terms of elasticity, performance and cost by using a workload scenario from a known dataset of taxis roaming in Beijing. We argue that the approach can be replicated in the design process of similar kinds of spatially distributed Internet-of-Things systems.

1 INTRODUCTION

The recent evolution towards an increasingly integrated world has at its basis novel types of pervasive systems achieved through new technologies and paradigms such as the Internet-of-Things (IoT). Such systems feature physically distributed devices and cloud computing infrastructure alike. This emergence comes along with new types of requirements and a need for increased assurances regarding the behavior of the overall physically-distributed systems [46], as they permeate more and more important aspects of human activity.

IoT systems which operate within a dynamic spatial environment are becoming ubiquitous; think of a taxi fleet within a smart city optimizing spatial distribution with respect to passenger demand, or a manufacturing floor co-habited by humans and robots. Those represent an important class of cyber-physical systems [8], which are faced with the manifold challenges that a dynamic spatial environment brings. They demand operational management to observe, evaluate and react to a constantly changing space. When the system is operational, analysis techniques situated at *runtime* are an essential prerequisite to ensure that possible changes occurring in the space –for example due to actions performed by active agents, or by the environment itself – do not lead to requirements violations. Typically, this can be achieved through a MAPE approach [41]; by (M)onitoring the spatial environment for changes, (A)nalyzing possible requirements violations, (P)lanning required countermeasures (e.g., moving a device from one point of space to another) and then (E)xecuting such actions and updating the shared model of space. We are not concerned

Authors' addresses: Christos Tsigkanos, Technische Universität Wien, Vienna, Austria; Martin Garriga, Universidad Nacional del Comahue, Neuquen, Patagonia, Argentina; Luciano Baresi, Politecnico di Milano, Milano, Italy; Carlo Ghezzi, Politecnico di Milano, Milano, Italy.

with choice and execution of appropriate counteraction measures that may be triggered to satisfy requirements, but with the requirements analysis activity itself. The rationale behind this is that verification of the system’s requirements –in the form of model checking– is (perhaps) the most computationally intensive activity, and should be appropriately and adequately supported by the system’s computational facilities. In contrast, counteractions can be actuated on the devices making up the IoT system directly.

Requirements within IoT systems often depend on the global state of the system (that is, the position of all devices in the spatial plane) and may be formally expressed into properties and verified [53]. However, providing assurances at runtime over the global state, e.g., through model checking, is computationally-intensive, making it unfeasible or impractical for the resource-constrained devices that populate the space in IoT scenarios [12]. A device may need to evaluate a property as part of the system’s business logic; a response with the property’s truth value may be required to, e.g., decide whether to perform some action or movement in the space. A naive approach would (i) place and (ii) populate instances of the model of space on each device and (iii) perform verification there – however, storage of the model within each device, communication of the global state on each device as well as the computational load that verification incurs on resource-constrained devices renders it impractical.

Since devices making up IoT systems are Internet-enabled, offloading the analysis computation to the cloud is a viable option [39; 57]. However, despite the general assumption of “infinite” resources on a cloud deployment, additional challenges arise regarding cost and resource allocation, specially when the workload is unknown at the system’s design time. In these cases, it may be difficult or even impossible to guarantee application’s Service Level Agreements (SLAs), e.g., on property evaluation response times. Besides, the cloud-based architecture should be elastic enough to handle such fluctuating workloads, while conciliating providers goals and client applications needs with efficient and scalable management of applications’ life-cycles. To this end, *microservices* [36] arise as the architectural weapon-of-choice, since they are small, modular, and independently deployable and scalable in an automated way. Additionally, with the rise of function-level compute instances through Function-as-a-Service (FaaS) models [38; 50], the fitness of generic cloud configurations needs to be re-evaluated for these applications.

Our approach lies within engineering of dependable systems operating in a discrete space arising from topological relations in the spatial environment, where the information abstraction of physical *location* or position of entities is inherently important [42; 52]. Change of such spatial position (mobility) in distributed IoT systems – e.g., for smart city, industrial or wireless sensor networks domains [46] – require to periodically re-evaluate the properties of the whole system. Firstly, assurance of such requirements may be critical, so techniques and methods for engineering dependable systems are applicable, such as formal verification. Secondly, spatial verification induces particular kinds of computational-intensive workloads, which should be supported by distributed software architectures. Finally, regarding system operation, leveraging microservices in a cloud infrastructure for spatial analysis while ensuring non-violations of SLA must be investigated.

As the cornerstone of our approach, we instantiate verification processes, integrating them within a cloud architecture for formally verifying requirements. Such requirements predicate on the topological distribution of devices comprising an IoT system at runtime, and formal verification is adopted to provide assurances of their satisfaction. Regarding system operation, we present the tradeoffs of the alternative deployments for such an architecture as microservices. The same type of tradeoff analysis and lessons learned can be applied to other systems in the class of systems discussed, and be easily extended to others. Specifically, the contributions of this paper are threefold.

- We address a challenging application scenario within formal assurance of global requirements of spatially-distributed IoT systems. We define the underlying model of space as a generic graph structure – a closure space [35] – and express global properties of the overall system in a topological spatial logic [26].
- We instantiate verification processes as (FaaS-ified [50]) microservices encapsulating spatial model checkers and instances of space models, integrating them to the service layer of the IoT-cloud architecture [46]. Finally, we position the microservices-based architecture into the IoT system environment, involving resource-constrained devices connected to the Internet.
- We propose alternative cloud deployments for such a model checking microservices-based distributed system, based on Virtual Machines (VMs), containers, and FaaS. We investigate their tradeoffs in terms of performance, scalability, elasticity and cost when evaluating global spatial properties at runtime. We further consider a novel *hybrid* deployment, combining VMs/Containers and FaaS. We discover whether the different cloud deployments are able to cope with given strict Service Level Agreements (SLAs) using a realistic workload scenario from a known dataset of taxis roaming in Beijing, as an representative instance of the class of systems we consider.

The rest of the paper is structured as follows. Section 2 presents a motivating example for spatial analysis at runtime. Section 3 details our proposal for reasoning on spatial properties of topological spaces. Section 4 discusses how to perform spatial verification at runtime as an IoT-cloud architecture, based on microservices. Section 5 details the different deployment alternatives and for such an architecture and their tradeoffs for offloading the reasoning machinery to the cloud. Section 6 discusses the evaluation of our approach and the alternative deployments through a well-known dataset of taxis in Beijing. Section 7 presents related work. Finally, Section 8 concludes the paper.

2 MOTIVATING EXAMPLE: BIKE-SHARING IN A SMART CITY

In this section, we introduce a motivating scenario as a running example of a spatially-distributed IoT system throughout the paper. We begin with a brief description of the scenario’s setting and then consider SLA goals that the system should exhibit as well as the challenges that arise.

Consider a bike-sharing scenario in a smart city, where bicycles roam in the physical space. As is common in such cases¹, bikes –or their users– maintain a connection to the network, rendering the system an instance of the IoT. A key issue in bike-sharing systems, is the concentration of bikes in certain parts of the city, other parts remaining empty, a phenomenon referred to as clamping. This may occur for instance around the start or end of the business day throughout the week, or due to sporadic events such as public concerts. The city administrator as the bike-sharing operator, has a significant interest to avoid this, as it leads to unavailability of bikes in other parts of the city when clamping occurs elsewhere. In our scenario, to mitigate this, the city identifies certain bike usage situations that describe optimal use of the bikes; a form of quality-of-service sub-goals that aim at avoiding undesired distribution of bikes. To make such cases attractive to bicycle commuters, the city offers rewards to bike-sharing users if a sub-goal is successfully satisfied for their bike. The sub-goals may reflect intuitive mitigations to the clamping problem, e.g. leaving the bikes around metro stations so that they can be readily re-used by others. However, clamping mitigations might be complex, capturing cases obtained from transportation analysis of the city. As an example, we consider the following two sub-goals –a simple and a complex one– of the bike-sharing system, concerning every single bike:

¹E.g., mobike: <https://mobike.com/global/>.

- (G1) Bikes should not be located close to each other at any point. Specifically, no bike should be located in a place nearby where another bike is.
- (G2) Bikes should be located in places in the city from which one can reach certain important landmarks –such as the main square– through the transportation network of bus and underground, but without traversing bridges. Additionally, no bus stops should be traversed if other bikes are located there.

Inherent in the formulation of the scenario above is a description of how various points in the city are connected, meaning that one can use a bike to go from a point to another. Such an *accessibility* model can be leveraged from existing map data or readily constructed by monitoring bike usage over time. The result is a graph structure, which has nodes corresponding to various points of the city and edges connecting those which are directly reachable – a topological map of the physical reachability of the city. For our example city, such a topological graph is depicted in Figure 1, where various city landmarks such as a park, a museum and transportation stops are connected due to being directly reachable from each other by bikes.

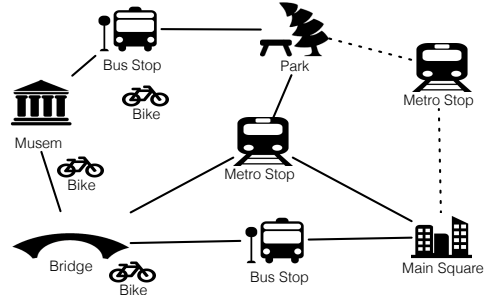


Fig. 1. Topological accessibility in a city.

Challenges inherent to the motivating scenario considered include both evaluating the truth values of sub-goals (G1) and (G2) for all bikes in the city at any moment, as well as architecturally supporting this analysis. Notice how the evaluation of sub-goals for each bike depends on the position of all other bikes in the city – the system’s global state– rendering the sub-goals *global spatial properties* [53]. Additionally, evaluation of such global properties is challenging because of computational and scalability concerns, requiring a dedicated architecture for the overall IoT system. A naïve approach would be to perform computation upon the devices constituting the IoT system themselves (assuming that they are communicated the global system state). In this case, the advantages are reduced latency and minimal persistent connectivity needs, as long as devices support the computation workload. However, the main disadvantage of this approach is the severe limitation on the computing resources of the IoT devices [49]. Even considering a top-of-the-range smartphone or tablet, SLA constraints and prohibitive battery drain make this approach unfeasible in practice [12] (see also Section 6 for workload comparisons).

3 REASONING ON SPATIALLY-DISTRIBUTED INTERNET-OF-THINGS SYSTEMS

In this section, we present our approach for reasoning about global spatial properties of (finite) models of topological space, captured by the spatial distribution of an IoT system. These models are closure spaces [35], a generalization of a standard *topological space*. Subsequently, we utilize a Spatial Logic for Closure Spaces (SLCS) [26], an extension of the topological semantics of modal logics to closure spaces, to support formal reasoning. In Section 3.1, we first provide a context for evaluating spatial predicates, by showing how a spatial evaluation model may be derived. This model is illustrated over the accessibility in the example city previously described. Given such an evaluation model, in Section 3.2 global properties can be verified.

3.1 Evaluation Model of Space

Our evaluation model is based on the established notion of trajectories [63] in geographical space, upon which many datasets and applications build. Such trajectories are often available as a domain model; for our running scenario, trajectories of bikes can be obtained e.g. by monitoring the bike-sharing system usage for a limited amount of time [7]. In the following, we demonstrate: (i) how a topological accessibility graph may be obtained from such a trajectory dataset; and (ii) how a discrete dataset consisting of traces of presences may be obtained. We obtain a topological *accessibility* graph in two steps:

- (1) Obtain the points-of-interest (POIs) in the city. These are nodes of the graph, each consists of a name, various attributes and geographical coordinates of the point in geographical space, sourced from some widely available repository (e.g. Open Street Map [1] or SPOI [21]).
- (2) Map trajectories of active entities over the POIs. If, based on a predefined distance, a trajectory passes through two points, these two points are considered *accessible* from one another, and an edge between the respective nodes is created on the graph.

Back to the running example, if a bike starts from the POI *museum* (Figure 1), goes through the *bridge* and arrives near the subsequent *metro stop*, then the three are considered (step-wise) connected, since accessibility was demonstrated by at least one trajectory. The resulting graph will have an edge linking the museum with the bridge and another linking the bridge to the metro stop. The predefined distance aims to mitigate sensing errors in the data, as well as resolution within the scale of which POIs are defined. The result of the process is a graph capturing accessibility of POIs in geographical space. Note how time/distance may be introduced to this model by time-stamping points of a trajectory when mapping over POIs and taking into account trajectory traversal time. The additional time dimension leads to a model describing trajectories as traces of presences – analogously, accounting distances enables quantitative reasoning.

The obtained accessibility graph bidirectionally connects POIs. This adjacency relation between POIs may induce a *closure space* [35], a mathematical model which can be used for formal reasoning with a spatial logic. A closure space is a notion originating from the field of mathematical topology, built upon what can be informally referred as the “least possible enlargement” of a set and exhibiting certain fundamental axioms. Formally, a closure space is a pair (X, C) where X is a set, and the closure operator $C : 2^X \rightarrow 2^X$ which assigns to each subset A of X its closure, such that for all $A, B \subseteq X$:

$$C(\emptyset) = \emptyset; A \subseteq C(A) \text{ and } C(A \cup B) = C(A) \cup C(B).$$

The elements of X are called the *points* of the closure space (X, C) . For any subset $A \subseteq X$, we define the *complement* of A in X as $\bar{A} = X \setminus A$. For a closure space (X, C) , for each $A \subseteq X$, the *interior* $I(A)$ of A is the set $C(\bar{A})$. Moreover, the *boundary* of a set informally refers to the set of elements which can be approached both from inside it and from outside of it. Formally, in a closure space (X, C) , the boundary of $A \subseteq X$ is defined as $\mathcal{B}(A) = C(A) \setminus I(A)$. Two more variants of boundary exist, the *interior boundary* $\mathcal{B}^-(A) = A \setminus I(A)$, and the *closure boundary* $\mathcal{B}^+(A) = C(A) \setminus A$. Finally, as shown in [26; 35], every graph for which the set of edges forms a binary relation induces a closure space, called a *quasi-discrete closure space*, by interpreting closure as the adjacency of

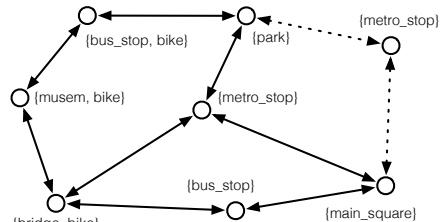


Fig. 2. Propositions associated to connected points form a closure model for the example city.

the nodes in our accessibility graph. Given a set of propositions P , a closure model [26] is a pair $M = (X, C), V$ consisting of a closure space (X, C) and a valuation $V : P \rightarrow 2^X$, assigning to each proposition the set of points where the proposition holds.

In Figure 2, the binary relation of *accessibility* between locations (POIs or landmarks) gives rise to a closure space. Thereupon, propositions associated with each point capture the POI name as well as presence of a bike.

3.2 Verification in Space

Building upon the previously defined fundamental operators of closure and boundary, we proceed to briefly outline the syntax and semantics of SLCS [26], a spatial logic for closure spaces. The logic will be evaluated upon models \mathcal{M} . The logic features boolean operators, a “one step” modality turning closure into a logical operator, and a spatially interpreted *surrounds* operator². Given that p ranges over a set of propositions P , the syntax of SLCS is defined by the following grammar:

$$\phi ::= p \mid \top \mid \neg\phi \mid \phi \wedge \psi \mid C\phi \mid \phi \mathcal{S}\psi. \quad (1)$$

In Formula 1, \top denotes true, \neg is negation, \wedge is conjunction, C is the closure operator, and \mathcal{S} is the spatial surrounds operator. Satisfaction $\mathcal{M}, x \models \phi$ of formula ϕ at point x in model $\mathcal{M} = ((V, C), v)$ is defined inductively on terms [25]. More complex logical operators can be defined based on the fundamental operators of closure and spatial surrounds. In the following, we recall fundamental ones [26; 53] being useful in the context of this paper and later used for the evaluation of our approach in Sec. 6. First, we consider the intuitive notion of *nearness*; points that are *near* another in a model can be found inside the set identified by (applications of) the closure operator upon this point.

Thus, nearness can be defined by nesting applications of the closure operator as $\mathcal{N}^n \phi \stackrel{def}{=} C^n \phi$. Operator \mathcal{N} can be applied arbitrarily often to predicate about points being in a defined *proximity* from each other. For example, based on the closure model of Figure 2, $\mathcal{N}(\{\text{museum}, \text{bike}\})$ will yield the three points on the leftmost part of the graph.

We further consider an operator \mathcal{T} based on the spatial surrounds operator [26] which captures a notion of *reachability*, defined as $\phi \mathcal{T} \psi \stackrel{def}{=} \phi \wedge (\neg((\neg\psi) \mathcal{S} (\neg(\phi \vee \psi))))$. $\phi \mathcal{T} \psi$ is satisfied for a point x if it satisfies ϕ and we can reach a point satisfying ψ while passing only by points satisfying ϕ . Then, based on \mathcal{T} , a more complex *reach through* operator \mathfrak{R} defined as $\phi \mathfrak{R}(\psi) \zeta \stackrel{def}{=} \phi \mathcal{T} ((\psi \mathcal{T} \zeta) \wedge (\psi \mathcal{T} \phi))$, is satisfied for a point x if x satisfies ϕ and there is a sequence of points starting from x , all satisfying ψ , reaching a target point satisfying ζ . The interested reader is referred to [26], as we consider defining complex operators as out of scope for this paper.

$$\text{bike } \mathfrak{R}(\text{!bridge} \vee (\text{bus_stop} \wedge \text{bike})) \text{ main_square}. \quad (2)$$

For an example illustrating the use of the derived \mathfrak{R} operator, consider the Formula 2 which specifies the `main_square` being reachable from a point which has the proposition `bike`.

4 SPATIAL ANALYSIS AT RUNTIME FOR THE INTERNET-OF-THINGS

In this section we discuss bringing spatial verification at runtime. After outlining essential computational properties of the model checking workload, we propose an application model where the workload is deployed on the cloud. Our proposal architecturally consists of software services and has as its main goals the satisfaction of SLAs and the optimization of the cloud resources used.

²Intuitively, closure and surrounds in a quasi-discrete closure space behave similarly to next and until in the temporal domain.

4.1 Spatial Verification as a Computation-Intensive Workload

Evaluation of spatial properties expressed in the SLCS syntax of Formula 1 over a closure model constitute a specific type of workload. The major computational component of this workload is the actual verification procedure – other minor tasks such as data structure manipulation are insignificant with respect to the computational burden of explicit-state model checking.

The model checking procedure evaluates SLCS properties over the closure model representing the current state of the system. Recall that the model consists of a graph structure (POIs and their adjacency relations, capturing accessibility in the model) as well as certain propositions that hold in points (e.g. presence of bikes). Following the model checking approach presented in [25], evaluation of a formula yields the set of points of the closure space where the formula is true. In our case, we return a truth value which indicates if the obtained set is non-empty. That is, the existence or absence of points in the closure model for which the evaluated formula is true. Thus, note that Formula 2 is violated for the bike in the museum in the upper left corner of Figure 2, but satisfied for the other two bikes located in the bus_stop and the bridge.

Given a spatially distributed system, representation of its discrete underlying structure as a set of points and their adjacency can be considered as virtually constant throughout its operation. With respect to the motivating example of Section 2, this corresponds to the topological structure of the city. The closure model additionally records the current positions of the bikes through propositions. Updating bike positions – changing the propositions assignment and subsequently verifying a property such as the one of Formula 2, is performed in a near-constant time. The actual verification burden induced depends on the size of the particular closure model, as well as on the complexity of the properties checked. However, the overall operation has challenging characteristics when considered as a computational workload. Spatial model checking as presented is a very computation-intensive operation, since it requires verifying the truth value of a complex formula. However, memory requirements are near-constant, since model checking requires keeping the closure model in memory, the size of which is known beforehand (i.e., it is the model size of the city). Such characteristics of the computational workload are exploited to design the microservices architecture presented in the following section, enabling spatial model checking to be performed at runtime while meeting strict SLAs.

Finally, we note that this renders the verification workload fundamentally different in contrast to model checking approaches on the cloud verifying temporal logic properties [14; 18; 44]. Given the above, we are not interested in partitioning or parallelizing verification computation, as our domain is spatial model checking, and the size of the space is virtually fixed per problem instantiation – we seek to evaluate properties that otherwise components in an IoT system would compute, given the particularities of the spatial verification workloads induced.

4.2 Service-based IoT Spatial Analysis in the Cloud

Given the limitations to perform computation in the devices themselves, offloading and deploying such computation to the cloud is a viable option [39; 57]. In the cloud, computing resources are typically provided through virtualization and containerization [10; 43], and there is an illusion of infinite resource availability thanks to horizontal scaling. However, accessing these resources may involve multiple hops of network communication, adding prohibitive latency in the processing of client requests. In this particular case of analysis however, computation is the major cause of delays. Thus, we propose an application model that brings together stateless components and immutable data, while stateful components are deployed separately. The proposed model prevents data consistency problems that would arise if stateful components such as databases were deployed in a distributed fashion. It also allows multiple service instances to coexist and be

deployed independently without the need for state migration, which is particularly important to cope with client mobility. Indeed, as devices will freely enter and exit geographical areas, even cloud providers may see considerable variations to their aggregate demand over time. In such a scenario it may be unfeasible to predict the origin and intensity of the workload [11].

When managing IoT services in the cloud, there are two conflicting goals: i) the satisfaction of SLAs, and ii) the optimization of the resources consumed by these services in the cloud. In this paper we focus in two kinds of application requirements, namely service latency and cost; simultaneously we target efficient and scalable usage of computational resources from cloud providers – mainly CPU, since our applications are computation intensive [10]. From the provider’s perspective, an efficient and scalable allocation of its (virtualized) resources must be allocated to cope with the SLA of each provided service. To be efficient, the allocation of resources should be able to mimic the corresponding fluctuations of demand, i.e., be highly responsive. Consequently, it is important that the mechanisms governing resource allocation be aware of the actual, and potential workload [11].

IoT system architectures bridge the gap between the physical and the virtual worlds, but entail multiple challenging factors. IoT architecture design involves networking, communication, and interoperability among heterogeneous devices [46], while the overall system must exhibit scalability. These devices may be heavily resource-constrained, e.g., in terms of computational power and battery constraints, while IoT applications have to meet stringent quality-of-service requirements, hindering the provision of computing and communication elements [60]. Due to the fact that things may move in the physical space and need to interact with others in real-time, an IoT architecture should be adaptive and decentralized, as well as support technology-agnostic event-based communication between heterogeneous devices serving as the fundamental architecture components. In this context, Service-Oriented Architectures (SOA) ensure interoperability among the heterogeneous devices [46] and allows to treat such a complex system as a set of well-defined simple services [9]. Due to these advantages, SOA has been widely applied as a mainstream architecture, for example in the context of Wireless Sensors Networks (WSNs) [23]. SOA applied to IoT provides extensibility, scalability, modularity and interoperability among heterogeneous things; in addition, the functionalities and capabilities are abstracted into a common set of services. Figure 3 provides a generic SOA view for IoT in the context of our problem, which consists of three layers [46]:

- Sensing/Actuating layer, integrated with available devices, responsible for obtaining their status and acting upon decisions based on analysis outcomes.
- Network layer, reflecting the infrastructure that supports wireless or wired connection among things, and
- Service layer, which creates and manages services used by end users and applications.

In our approach, IoT services are indeed microservices [36; 45], as they are small, modular, communicate with lightweight mechanisms (often through an HTTP RESTful API) and are independently deployable by fully automated machinery. Back to our running example, a bike’s location sensors (sensing layer) emit heartbeats with its coordinates, either periodically or upon certain events (e.g. locking/unlocking, or in proximity to certain landmarks or POIs). Such heartbeats should trigger the evaluation of the global model in order to determine whether the defined sub-goals hold for the current position of the bike (as well as the positions of all other bikes). The event travels across the network infrastructure towards the cloud (Network layer), reaching the cloud service in charge of processing it (Service Layer). Finally, the result of the evaluation travels back to the actuator layer, to undertake the corresponding actions e.g. through the mobile device of the user (e.g., providing feedback regarding penalties or rewards).

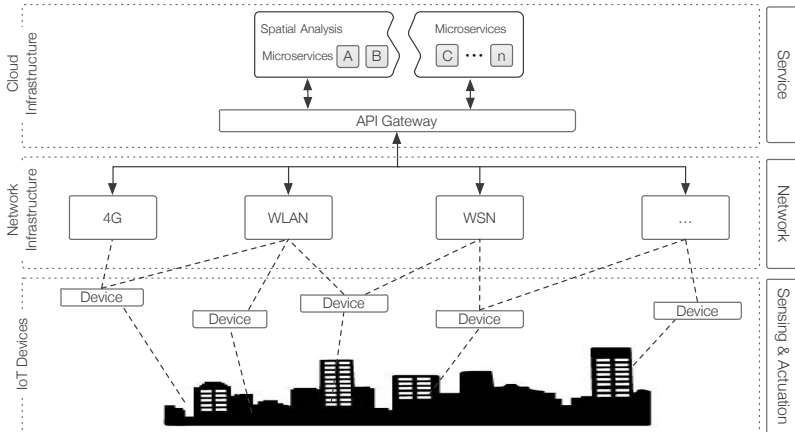


Fig. 3. IoT and microservices architecture for materializing spatial reasoning on the cloud.

5 CLOUD DEPLOYMENTS FOR IOT SPATIAL REASONING

This section details the materialization of our proposal as a microservices architecture for IoT (based on the SoA-IoT interplay discussed in Section 4.2), which allows to offload the reasoning on space presented in Section 3. We instantiate verification processes integrating them to the service layer of the IoT-cloud architecture [46] of a spatially-distributed IoT system involving resource-constrained devices connected to the Internet. Applications with specific processing demands, including the computation-intensive ones discussed through this paper, have historically required special configurations such as compute- or memory-optimised virtual machine instances [50]. With the recent rise of function-level compute instances through Function-as-a-Service (FaaS) models, the fitness of cloud configurations needs to be re-evaluated for these applications. Subsequently, we propose alternative cloud deployments for such a model checking microservices-based architecture, based on Virtual Machines (VMs), containers, and FaaS, as well as a *hybrid* deployment setup.

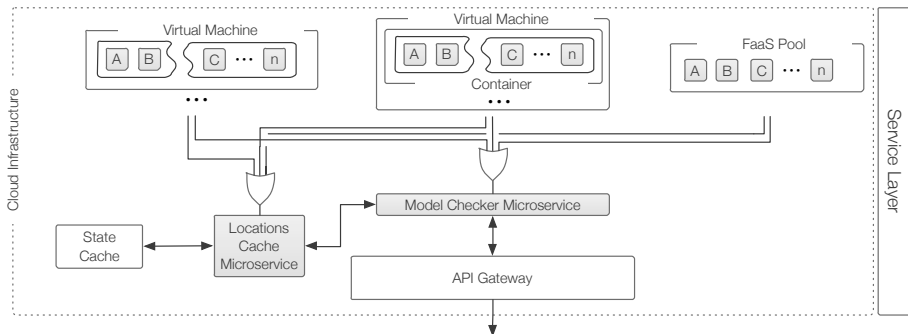


Fig. 4. Deployment diagram for the service layer.

The service-based architecture consists of two basic building blocks, materialized as two separate microservices³ Stateless components reflecting solely computation (i.e., the model checking procedure evaluating spatial properties of Section 3.2) and mostly immutable data (i.e., the topological graph of space of Section 3.1) form the model checker microservice, which receives the current system state and a property, and verifies if such property holds. The system state refers to the current spatial distribution of entities in the city as a closure space, while the property captures the QoS requirements that should be evaluated upon the current state of the system. Stateful components form the location-cache microservice, which provides an interface to the current

³Complete source code as well as accompanying evaluation models are available at [2].

global state of the space, i.e., the location of the devices at a given point in time. Needless to say, both microservices provide HTTP/RESTful interfaces. In this manner, those microservices can be reused, scaled and maintained individually; therefore, the software and hardware components in the IoT deployment can be reused and upgraded efficiently. In the following, we propose concrete deployments and instantiations of those microservices with the goal of enabling computation offloading to the cloud, and then discuss their advantages and disadvantages. Evaluation of the different deployments upon a real-world workload is provided in Section 6.

5.1 Elementary Cloud Deployments

Monolithic Deployment. The whole logic of the application (i.e., the two microservices and the database) lives within a single deployment unit. This is the straightforward option for the cloud as it can be deployed in a single monolithic server, typically a virtual machine (VM), which substantially reduces the operational burden [34]. However, scalability is constrained in two ways: first, given the limited computing resources of such a monolithic deployment, and second, because the scale unit is the whole application, even though different microservices have unequal loads and may need to scale separately. Alternatively, the monolith can be deployed in a multi-server environment, with a load balancer to distribute the load among multiple VMs. However, adding VMs to handle workload bursts may take several minutes [10], which can cause severe SLA violations. Again, the fact that all components scale together (despite their actual utilization) multiplies the resource overprovisioning, something not cost-effective.

Containerized Microservices Deployment. Each microservice may be developed separately using different technological stacks, deployed and scaled independently in a containerized manner. Containers are executed directly on top of the host operating system (typically a VM), optionally with the help of a container manager (e.g., Docker [5]). Each container typically hosts one microservice, with its own platform stack and application code. Containers have various advantages when compared to VMs: they are more lightweight, and faster to boot and terminate [33]. This allows for faster and more reactive scalability: if scaling VMs takes several minutes, adding a container or changing its resources takes seconds. Since multiple applications that share the same virtualized resources may see different workloads, the finer granularity and the higher adaptation speed of working with containers allow us to use resources more efficiently [10]. The disadvantages of this approach may rely on the increased operational burden (since there are several deployment units to manage and scale) and, still, SLA violations that may occur during scaling-up actions.

Functions-as-a-Service Deployment. FaaS, also known as serverless [38; 48; 50]) has been recently proposed as an alternative cloud paradigm in which business functionality is provided without pre-allocating computational resources. Instead, shared resources (e.g. containers) are used to provision and execute functions on demand, typically in a few milliseconds. Therefore, microservices are implemented as functions (often called *lambdas* [3]) allowing a straightforward deployment to the cloud. The horizontal scaling in FaaS is completely automatic, elastic, and more reactive than the typical solutions of scaling a virtual machine or spinning up containers against bursts of workload [38]. A disadvantage of this approach may be vendor lock-in, since specific cloud solutions (e.g. AWS Lambda[3], IBM/Apache Openwhisk[6], Azure Functions[4]) are tied to the service ecosystem from the same cloud provider. Additionally, functions are stateless by definition, which limits the applicability of the model – stateful microservices should store the state outside of the function itself, as done for the location-cache in our deployment. Since the computationally-intensive model checking workloads we investigate require minimal state, a FaaS deployment fits such workload requirements particularly well.

5.2 Spatial Analysis Workloads over a Hybrid Deployment

A deployment based on VMs or containers, while reducing operational burden may fail to perform with respect to SLAs due to slow scaling actions. On the other hand, the elasticity that FaaS offers should be leveraged. To this end, we propose a novel *hybrid* deployment: a baseline infrastructure based on VMs or containers to serve normal, expected workloads, plus a FaaS fallback infrastructure to serve unexpected workload peaks. The rationale is to leverage lower cost of usual scaling options while fluctuations are managed by highly elastic and reactive FaaS components. Figure 5 shows our proof-of-concept hybrid deployment featuring a Monolithic server plus a FaaS pool.

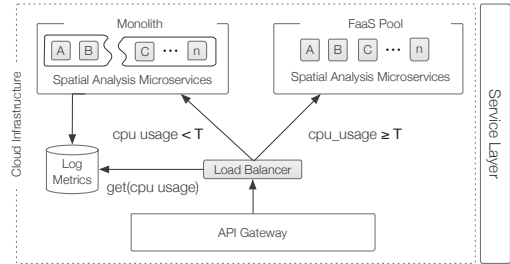


Fig. 5. Sample Hybrid deployment featuring FaaS and a monolith server behind a load balancer.

Instead of invoking directly microservices, a Load Balancer is in charge of forwarding the request to the corresponding server or function, depending whether the CPU load of the server is higher than a given threshold T . The request is then forwarded to a lambda function when the server is overloaded and thus may fail to process the request. In the following section, we evaluate this the alternative deployments with a sample workload scenario.

6 EVALUATION

For evaluating the proposed approach and deployments, we developed tool support and a proof-of-concept implementation based on the SLCS topochecker [26], available online along with accompanying material [2]. We deployed the prototype to the different cloud alternatives discussed in Section 5. Our evaluation goals target realization and applicability of our approach for spatial service-based analysis for IoT applications. Concretely, we aim to:

- Investigate feasibility of the approach through concrete deployments in the cloud;
- Compare different deployments in terms of latency and SLA in the context of a normal workload, i.e., one known beforehand, to provide a baseline for the required computational resources and cost on each alternative, and
- Perform stress tests over the different deployments by increasing the workload rate, to assess scalability of each alternative.

Our evaluation model and workload is sourced from the Microsoft T-Drive trajectory dataset [61; 62], which contains trajectories of 10,357 taxis in Beijing spanning one week. The number of coordinate points and total trajectories distances in the dataset is in the order of millions. Based on the T-Drive dataset, we derive a closure model as per Section 3.1. We source POIs of Beijing from the OpenStreetMap repository [1]. To obtain discrete trajectory data from the trajectory dataset, we record presences of entities (taxis) over time and space – i.e., over the points of the accessibility graph previously defined. Specifically, for every single taxi in the dataset, its trajectory – as a line – is placed over the POIs of Beijing. If the distance from a point in the trajectory to a POI is less than a predefined length (10 meters for our evaluation purposes), a time-stamped presence of the taxi in that POI is recorded. The result is a discrete, timestamped sequence of presences of entities in POIs, which ultimately leads to a graph capturing accessibility of such POIs in geographical space. For example, if a taxi starts from a POI "Tsinghua University", goes through a bus stop and arrives at "Peking University", then the three are considered (step-wise) connected, since accessibility was demonstrated by at least one trajectory. The predefined distance aims to mitigate

sensing errors in the data, as well as resolution within the scale of which points-of-interest are defined. For the T-Drive dataset, an example trace would record $\{taxiA, date, TsinghuaUniversity\} \cdot \{taxiA, date, BusStopA\} \cdot \{taxiA, date, PekingUniversity\}$. Evaluation datasets and derived models are available in accompanying material [2].

We present traditional deployments setup on Section 6.1, and the experimental results obtained in Section 6.2. Then, we evaluate the hybrid deployment alternative in Section 5.2, and subsequently discuss threats to validity.

6.1 Experiments Setup

Given a closure model and a trace of presences as previously defined, we proceed to consider the workload induced by the active entities in the space – in this case the Beijing taxis – given the scenario presented. We assume that each registered presence of a taxi in a POI, also implies the evaluation of certain spatial properties (similarly to Section 2) by calling the model checker microservice. The evaluation of the property is based on the global state of *all* taxis in the Beijing dataset at that time.

We limit our evaluation clock-time to a window of 24 hours in Beijing. Figure 6 depicts single requests (i.e. taxi presence heartbeats) over time for a representative weekday, where one can observe that the workload exhibits a sine wave following the working hours. We consider a closure model of 36805 edges, with 15456 propositions over 5152 nodes [2].

The performed experiments were threefold. We defined an SLA of 30 seconds for the response time of the properties evaluation service. For the first experiment, we executed the sample window of 24 hours of Figure 6 and profiled the necessary infrastructure (for each deployment) in order to satisfy the SLA within the whole period. This allowed us to determine the baseline infrastructure and cost to serve a typical workload based on the evaluation model. Then, in the second experiment we proceeded

with a stress test, by reproducing one hour of workload with different time multipliers, i.e., simulating that the events arrive with a faster rate. This allowed us to assess the scalability of each deployment, in a scenario in which the workload rate is increased with time multipliers. Finally, the third and last experiment assesses the performance of our hybrid approach with respect to using only VMs or FaaS. The hybrid combines the best of the previous alternatives: a baseline infrastructure with preallocated resources (in the form of VMs or Containers) to serve normal workloads, plus a FaaS fallback infrastructure to serve unexpected peaks in the workload.

The example properties to be checked (Formulae 3-5) are system-wide goals in line to the scenario of Section 2, that aim to assign penalties or rewards to taxis in order to avoid clamping in the city. Specifically, P1 encodes the requirement that a taxi shall be able to reach department stores traversing points of subway or bus stops. However, traversal may not go through a hospital. P2 encodes that the taxi is always in the vicinity of a hospital or a hotel (by two steps away), but not immediately next to restaurants. P3 specifies that a taxi should be within an area characterized by proximity to tourist attractions (e.g. the Forbidden Palace) and subways, up to zoo’s (e.g. the

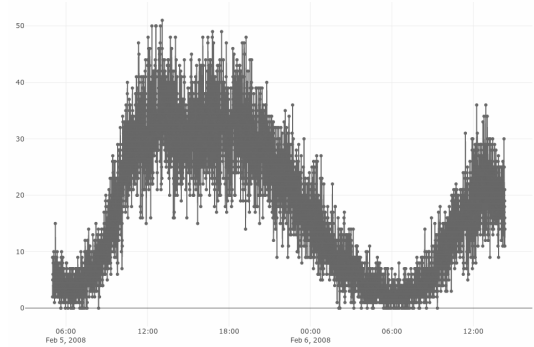


Fig. 6. Beijing traces [62] dataset evaluation workload [2] as taxi presences (y-axis) near points-of-interest [1] over the course of hours of the days of Feb 5-6 2008 (x-axis).

Beijing Zoo). Such a property, utilizing the \mathcal{T} operator of Section 3.2 could be used to reward taxis implementing a temporary city-wide policy of attracting tourists which are located in the center of the city to go to the zoo, since they will be able to find more taxis within this area. Note that specification formulae utilize the complex operators previously defined – the corresponding elementary formulae are sizable. The propositions used in the specifications are found in the generated dataset available in accompanying material [2].

$$(P1) \text{taxi} \ \mathfrak{R} \left(\left(!\text{HEALTHHOSPITAL} \vee \left((\text{TRANSPORTSUBWAY} \vee \text{TRANSPORTBUSSTOP}) \wedge \text{taxi} \right) \right) \text{DEPARTMENTSTORE} \right) \quad (3)$$

$$(P2) \text{taxi} \wedge \mathcal{N}_2(\text{ACCOMMOHOTEL} \vee \text{HEALTHHOSPITAL}) \wedge \mathcal{N}(\text{FOODRESTAURANT}). \quad (4)$$

$$(P3) \text{taxi} \wedge \mathcal{N}_2(\text{TOURISTATTRACTION} \wedge \text{TRANSPORTSUBWAY}) \ \mathcal{T} \ \text{TOURISTZOO}. \quad (5)$$

We deployed the microservice architecture described in Section 4.2 (Figure 4) in four different environments. We note that the workload, being a model checking procedure, is heavily computation-intensive, but its memory requirements are minimal, so we ignore RAM specifications of the infrastructures for clarity.

Device: Local execution of the model checking load in an IoT resource-constrained device (e.g. a mobile phone), featuring an ARMv6 A53 1GHz CPU.

Monolith: Large VM instance containing the two microservices in a single deployment package. The VM ranges from 16 CPUs (c5.4xlarge AWS instance) for light workloads, up to 64 CPUs (c5.18xlarge) for heavy workloads.

Containers: AWS Elastic Container Services (ECS⁴), where each microservice is containerized, deployed and scaled independently. The infrastructure ranges from 10 to 25 instances with 2 CPUs each (t2.medium), depending on the workload.

Lambdas: Functions-as-a-Service as AWS Lambda functions were configured with 2 vCPUs each⁵, scaling automatically from zero to the level of parallelism required, according to the workload.

The services deployed correspond to a location-cache microservice as an in-memory key-value store, while the checker microservice as an encapsulation of the model checker [25], which receives the current state and a property, and verifies if the property holds. The IoT device, featuring similar processing power to a mobile phone, encapsulates a model checker procedure compiled accordingly per its software architecture stack. The relevant implementations are available in [2].

6.2 Cost and Scalability Evaluation Results

Table 1 shows the infrastructure and cost results for the first experiment, considering 24 hours of execution (resulting in 8864 model checker invocations) sliced in periods of three hours. The third column (*Cores*) shows the approximate number of cores needed to serve such a workload within the SLA of 30 seconds per call. Then, each deployment is scaled to be close to that number during each period.

Results show that, when the workload is known beforehand, the Containerized Microservices solution is the most cost effective: \$456.96/month, considering a month as 21.73 working days (the workload considered decreases significantly during weekends). Subsequently, the FaaS deployment costs \$771.67/month, calculated upon the amount of memory assigned to the function, and the exact number of invocations and execution time (12 seconds on average), which eases its fine-tuning. The less cost effective solution is the Monolithic Server deployment, at a cost of \$932.63/month. This is due to the fact that the Monolithic Server incurs resource over-provisioning most of the time, even

⁴<https://aws.amazon.com/ecs/>

⁵CPUs are assigned proportionally to the RAM (2GB in this case) – <https://aws.amazon.com/lambda/>

Table 1. Baseline infrastructure and cost for the cloud deployments considering 24h workload.

Hour	Calls	Cores (aprox)	Containerized Microservices		Monolithic Server				FaaS Deployment	
			VMs*	cost	VM	cores	memory	cost	sec/call	cost
6-9	312	20	10	1.4	c5.4xl	16	32	2.04	12	1.25
9-12	1283	40	20	2.8	c5.9xl	32	64	4.6	12	5.13
12-15	1663	50	25	3.5	c5.18xl	64	128	9.18	12	6.65
15-18	1676	50	25	3.5	c5.18xl	64	128	9.18	12	6.70
18-21	1668	50	25	3.5	c5.18xl	64	128	9.18	12	6.67
21-00	1202	40	20	2.8	c5.9xl	32	64	4.6	12	4.80
00-03	705	30	15	2.1	c5.4xl	16	32	2.04	12	2.82
03-06	355	20	10	1.4	c5.4xl	16	13	2.04	12	1.42
Day	8864			21				42.86		35.46
Month**	192614			456.96				932.63		771.67

*t2.medium AWS instance, 2 CPUs, 4GB memory

**Month = 21.73 working days

though we are adjusting the size of the VM for each period. Finally, we also profiled the execution in the mobile device, which represents a single taxi in the dataset. Rather than testing with the workload (multiple taxis), we performed sequential calls. This captures the case where instead of a cloud solution, one chooses a deployment where properties are evaluated on the devices making up the IoT system. Execution results shown that single invocations take 99 seconds on average to complete, which violates our SLA of 30 seconds, making it unfeasible in practice (barring e.g. battery drain effects). Nevertheless, we additionally included the mobile device in the subsequent experiment for completeness.

For the second experiment we considered a representative lapse of one hour of workload, from 11am to 12am of February 5 (Figure 6), which is the beginning of a daily peak comprising 536 data points. We executed stress tests over the different Cloud deployments, by multiplying the rate in which requests arrive (one per each data point). Each request triggers one invocation of the model checking procedure. We stretched the 536 invocations in 12 minutes (5x faster rate), 6 minutes (10x), 3 minutes (20x), 1.5 minutes (40x) and 1 minute (60x). Then, we measured the total time per invocation, considering the SLA of 30 seconds defined before, and disaggregated it by computation time and wait time, i.e., when the request is queued waiting for computation resources.

Figure 7a depicts the median time per invocation for the different rates and deployments. Additionally, we plotted the mobile device (rightmost column), which executes sequential requests independently of the rate: as soon as one request is processed, another one is fired. The Containerized Microservices in the ECS cluster presented the worst performance overall, with requests timing out (error 504) with 20x rate or more, which explains the absence of the yellow column for 40x and 60x rates. Since a scaling-up action (that can demand several seconds [10]) triggers when more containers are needed to handle the workload, requests timeout in the meantime. The FaaS alternative held an (almost) constant performance with 12 seconds per invocation for all workloads. Even though commercial FaaS solutions can give the impression of "infinite" scalability, the tradeoff between resource availability and cost (Table 1) has to be carefully considered. The Monolithic deployment performed better (between 9 and 12 seconds per invocation) up to 20x rate, then started to decrease its performance due to resource unavailability. As discussed in Section 5.1, the monolith could also scale in a multi-server environment, although adding VMs to handle workload

bursts as in this experiment may take several minutes [10], which can cause severe SLA violations in the meantime. Finally, the Mobile Device computes the requests sequentially with a constant performance of an average of 99 seconds per invocation. Needless to say, such processing time is prohibitive as it violates the defined SLA. Figures 7b and 7c show the actual computation and waiting time respectively. Recall that waiting occurs when no computation resources are available at the moment the service is invoked. One can observe that for all the deployments, computation time remains almost constant, the underlying issue being the wait time which substantially increases for faster workload rates. In such cases, requests are queued waiting for available computational resources, which decreases the overall system performance.

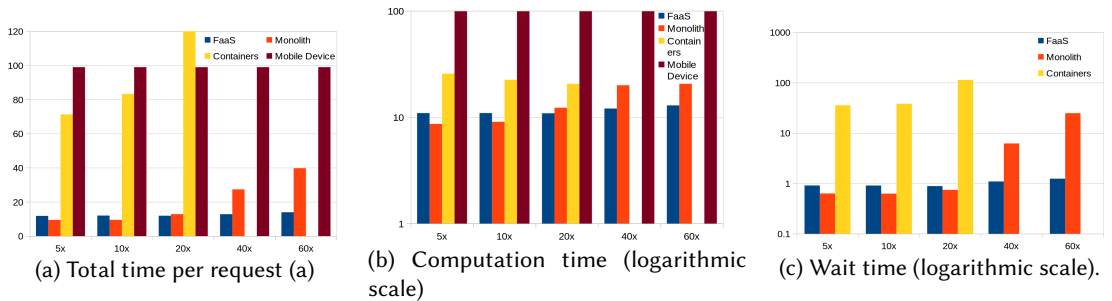


Fig. 7. Average time per request for different deployments and workflow rates.

From the two first experiments we can highlight the following. First, the solution based on the mobile device, even though it processes only its own requests, has a prohibitive computation time of 99 seconds on average, making it unsuitable in practice. The unfeasibility of performing computation locally calls for computation offloading to the cloud, with all its different flavors. The two solutions that performed better were FaaS and Monolithic. The FaaS deployment using AWS Lambda held a constant performance even under dense workloads, being more reactive and scalable than the Containerized and Monolithic alternatives. The option of the Monolithic deployment in traditional AWS VMs, performed well for medium workloads, but not against peaks or faster rates, specially when the latter are unknown beforehand. Finally, the Containerized Microservices deployment, materialized through an AWS ECS Cluster did not seem suitable for this scenario, mainly because of scaling-up delays (even when containers take only seconds to bootstrap). Regarding the overall cost, when the workload is more or less known (e.g., a normal workday in our taxis scenario) FaaS become more costly in comparison to other deployments that allow one to preallocate resources in advance. This opens the challenge to consider the hybrid approach, as follows.

Hybrid vs. monolithic and FaaS deployment. For the third experiment, we implemented the hybrid deployment of Section 5.2 (Figure 5) and assessed its performance with respect to other cloud deployments – namely monolith and FaaS. Our proof-of-concept hybrid deployment features a Monolithic server plus a FaaS pool, transparent for the end user thanks to the API Gateway, which provides a single REST endpoint to reach any of the checker microservices. Instead of calling directly the microservice, the API Gateway fires a lambda function which acts as the Load Balancer that forwards the request to the corresponding server/function, depending on the load of the server being higher than the configurable threshold T .

Figure 8 depicts the total, computation and wait time per invocation for the different request rates, considering the hybrid deployment, plus lambda and monolith as the two best alternatives according to the previous experiments (see Section 6.2). Note that the plotted values correspond to 20x rates (3 requests/second) and higher, since for lighter workloads the hybrid deployment is not necessary (100% requests being forwarded to the monolithic server). For 40x rates, the monolith served 72% of requests and FaaS 28%. For 50x rates, 67% and 33% respectively, and for 60x the

distribution was 65% and 35%. For 20x rates all deployments performed equally. For faster rates, the monolith server exponentially decreases its performance mainly due to wait times (insufficient computational resources). Meanwhile, lambda and hybrid held an almost constant performance of 12 and 20 seconds per call respectively.

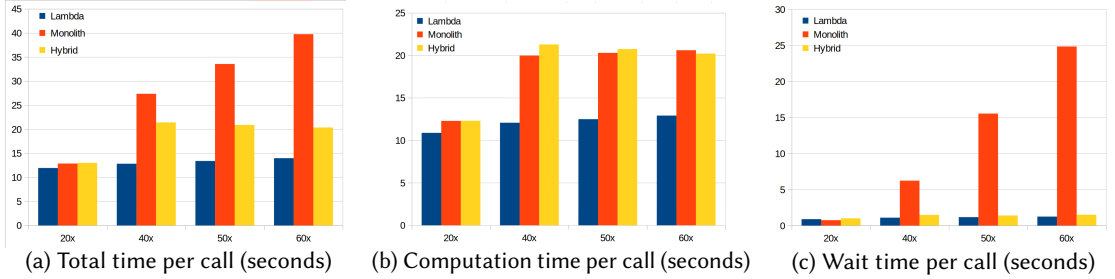


Fig. 8. Time (y-axis) comparison for different deployments and workflow rates (x-axis).

Figure 9 shows the maximum, minimum and median times per call for FaaS, Monolith and Hybrid deployment, regarding the 30 seconds SLA previously defined. FaaS holds a constant performance without incurring SLA violations. The Monolithic alternative incurs SLA violations for 40x rates or higher, not only for the maximum but also for the median times. Meanwhile, the hybrid deployment shows SLA violations only for the maximum times, but not for the median. This is due to the requests that arrive at the beginning of a peak, that is, when the server starts to be overloaded, but the LB has not yet detected it. Conclusively, the hybrid deployment combines the best of the two worlds: preallocated resources to serve normal workloads, while lambda functions allow to serve unexpected peaks on the workload in a reactive and cost-effective fashion. In all cases, the lambda load balancer introduces little overload, less than 300ms per request on average.

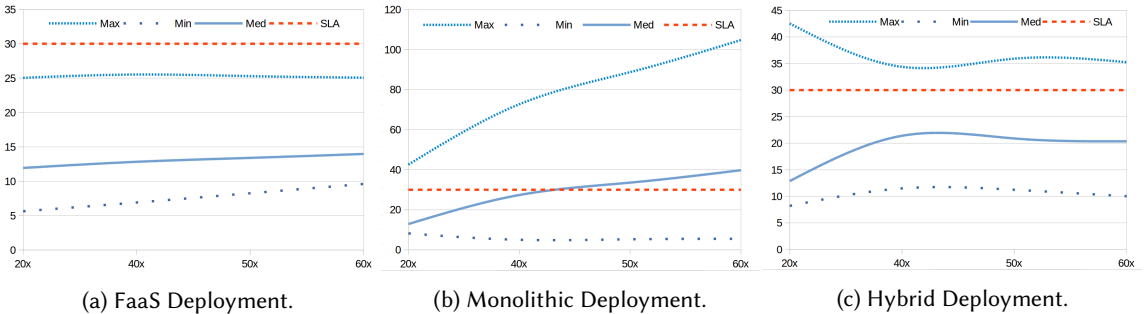


Fig. 9. Total times (max, min, median) and SLA violations for FaaS, Monolithic and hybrid deployments.

Finally, in terms of threats to validity is it worth mentioning that the experiments targeted only one sample model checker [26], and three spatial properties. Further tests may be needed, considering other properties and implementations that may affect processing times. Additionally, one may argue that the defined SLA of 30 seconds is arbitrary. This should be considered only as a reference value in the context of our experiments, whilst an actual SLA may be even lower. In those cases, the model checker implementation and deployment configuration (e.g., timeouts, location cache lifetime, data locality) should be optimized to meet the requirements of a real-time application. Regarding the workload behavior, the cloud deployment alternatives that perform resource preallocation (i.e., monolith server, containers) require to know the workload beforehand. When this is not possible, it may be better to rely on more flexible deployments such as FaaS and hybrid (the latter with a minimal amount of preallocated resources), which can prove more

expensive in the long term. It should be also taken into account that the hybrid deployment used for the third experiment constitutes a proof-of-concept. Fine-tuning the deployment can result in further improvements, in particular regarding the threshold value T for CPU usage, and caching policies for the load balancer. Last but not least, using containers and/or edge servers for deploying the hybrid alternative could substantially improve its performance [10–12], these options being considered as a significant avenue of future work.

7 RELATED WORK

We presented an approach for model checking requirements predicating on the spatial distribution of devices comprising an internet-of-things system at runtime, leveraging a cloud architecture based on microservices. Consequently, we classify related work into three categories. First, we review foundational work on spatial reasoning, positioning our work. Then, we discuss relevant approaches that have considered model checking and computational intensive tasks using cloud infrastructures. Later, we discuss related engineering approaches that aim at systems utilizing spatial verification for various purposes, in context of which our technical framework may be used.

Topological relations have been traditionally considered in the context of database systems, query languages [40] and logics for spatial data analysis in geographical information systems [15], focusing on the elements of a geometric model [30; 31]. In addition, spatial logics have also been studied in the context of process calculi [20], where the typical theme is predication against the structure of agents, also with applications to graph databases [19]. Our choice of a spatial logic over closure spaces [25] as a basis for our spatial reasoning support is that closure spaces are a generic mathematical concept serving as the interface between arbitrary binary relations. We further note that quantitative reasoning in our application domain appears promising [17], something we identify as future work.

The recent extensive availability of cloud processing platforms has fueled notable approaches for classical model checking [14; 44], proposing a technological transition to exploit the new available architectures. Approaches target classical, foundational verification of temporal logics such as CTL and LTL [29]. The prevalence and availability of the cloud has created an increasing interest in parallelizing and distributing verification techniques. Generally, the exploitation of cloud resources aims at increasing the memory available and reducing the overall time required by verification by employing distributed techniques [13; 32] or by splitting the given state space into several partial state spaces [18]. We note that the model checking processes that we instantiate, do not partition or parallelize verification computation, as our domain is different – we seek to evaluate properties that otherwise devices in an IoT would compute, given the particularities of spatial verification. To the best of our knowledge, spatial model checking within our cloud context has not been considered before; spatial verification exhibits particular kinds of computational workloads, as we observed in Section 4.1. Moreover, the potential offered to engineer applications in novel domains such as the IoT, has not been exploited, as discussed in the following.

We position our approach within engineering of dependable systems operating in a discrete space arising from topological relations in the spatial environment, where the information abstraction of physical *location* or position of entities is inherently important [42; 52]. Topological relations have found application in various domains, where formal reasoning has proved beneficial, such as in foundational approaches [26; 28]. In previous work [53] we considered spatio-temporal model checking of evolving spaces and advocated that the topology can provide a system with awareness of multiple characteristics [52; 54; 55]. Closure space-based reasoning with the SLCS logic has been considered for data correctness in vehicle location data, in the context of collective adaptive systems [24]. Moreover, a combination of CTL and SLCS is developed [27] to study bike sharing systems. In [51], a combination of metric and spatial logics has been proposed for

verification of safety properties in cyber-physical systems, while run time verification of spatio-temporal behaviours of complex systems is studied in [47], extending Signal Spatio-Temporal Logic with SLCS. For our evaluation purposes, we utilize a trajectory dataset, so we succinctly discuss relevant approaches. Trajectory-based reasoning has been studied extensively in scientific literature, where trajectories are typically modeled as streams of spatio-temporal points. Recently the focus has been on the use of semantic abstractions of raw mobility data, including not only geometric information but also knowledge extracted jointly from the mobility data and the application domains information [59]. We note advanced data mining techniques which tackle related problems of low sampling-rate and uncertain features of trajectory data [22], their context [56], or forecasting [58].

8 CONCLUSIONS

In this paper, motivated by the need for requirements assurance of spatially-distributed IoT systems, we proposed a technical framework for their spatial formal verification at runtime, leveraging a microservices architecture for the cloud. We defined the underlying model of space as a topological graph structure paving the way for expressing and verifying spatial logic properties. We devised a microservices architecture where we instantiated spatial model checking processes in a service layer, and proposed alternative cloud deployments; Virtual Machines, containers and the recent Functions-as-a-Service. After investigating their tradeoffs when evaluating global spatial properties at runtime, we proposed a novel hybrid deployment, where preallocated resources are augmented with an additional FaaS fallback infrastructure serving unexpected peaks in the workload. Our evaluation with a realistic workload scenario investigates feasibility, performance, scalability and elasticity of different deployments.

Regarding future work, we aim to first challenge analysis assumptions of IoT-cloud coupling as well as enable the specification and verification of more complex spatial properties. Today's IoT-cloud architectures assume that cloud resources are always available, and where software services then provide the analyses, business logic and data for devices to take action. However, novel requirements driven by technological advancements such as timeliness, locality of computation or privacy may dictate analysis to take place closer to IoT devices, at the *edge* of the network [37]. For example, the round-trip network latency for spatial analysis (for smaller spatial models) may be prohibitive when analysis is performed on the cloud [12], or fault-tolerance requirements may suggest that the central point of failure that the IoT-cloud coupling exhibits is not acceptable. As such, we aim to support analyses on the *edge* [11] where the challenge is decentralizing verification as well as the appropriate data to support it, putting more focus on edge devices [16] such as IoT gateways, network devices, cloudlets, and small clouds. On the analysis side, we have demonstrated that relevant qualitative spatial requirements may be expressed with the SLCS logic. We plan to extend this predication to include quantitative [47] and metric [51] aspects.

REFERENCES

- [1] 2017. Open Street Map. <http://wiki.openstreetmap.org/>. (2017).
- [2] 2018. Accompanying datasets, models and open source implementation. <http://178.62.206.217/lambda-topochecker/>. (2018).
- [3] 2018. AWS Lambda. <https://aws.amazon.com/lambda/>. (2018).
- [4] 2018. Azure Functions. <https://azure.microsoft.com/en-us/services/functions/>. (2018).
- [5] 2018. Docker. <https://www.docker.com>. (2018).
- [6] 2018. IBM/Apache Openwhisk. <https://openwhisk.apache.org/>. (2018).
- [7] Natalia Adrienko and Gennady Adrienko. 2011. Spatial generalization and aggregation of massive movement data. *IEEE Transactions on visualization and computer graphics* 17, 2 (2011), 205–219.
- [8] Radhakisan Baheti and Helen Gill. 2011. Cyber-physical systems. *The impact of control technology* 12 (2011), 161–166.
- [9] Luciano Baresi, Martin Garriga, and Alan de De Renzis. 2017. Microservices Identification Through Interface Analysis. In *European Conference on Service-Oriented and Cloud Computing*. Springer, 19–33.

- [10] Luciano Baresi, Sam Guinea, Alberto Leva, and Giovanni Quattrocchi. 2016. A discrete-time feedback controller for containerized cloud applications. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ACM, 217–228.
- [11] Luciano Baresi, Danilo Mendonca, Martin Garriga, Sam Guinea, and Giovanni Quattrocchi. 2018. A Unified Model for the Mobile-Edge-Cloud Continuum. *ACM Transactions on Internet Technology* (2018). In Press.
- [12] Luciano Baresi, Danilo Filgueira Mendonça, and Martin Garriga. 2017. Empowering Low-Latency Applications Through a Serverless Edge Computing Architecture. In *European Conference on Service-Oriented and Cloud Computing*. Springer, 196–210.
- [13] Jiri Barnat, Lubos Brim, and Jakub Chaloupka. 2003. Parallel breadth-first search LTL model-checking. In *Automated Software Engineering, 2003. Proceedings. 18th IEEE International Conference on*. IEEE, 106–115.
- [14] Carlo Bellettini, Matteo Camilli, Lorenzo Capra, and Mattia Monga. 2013. Distributed CTL model checking in the cloud. *arXiv preprint arXiv:1310.6670* (2013).
- [15] Roger S Bivand, Edzer Pebesma, and Virgilio Gómez-Rubio. 2013. *Spatial Data Import and Export*. Springer.
- [16] Flavio Bonomi, Rodolfo Milito, Jiang Zhu, and Sateesh Addepalli. 2012. Fog computing and its role in the internet of things. In *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*. ACM, 13–16.
- [17] Luca Bortolussi and Laura Nenzi. 2014. Specifying and monitoring properties of stochastic spatio-temporal systems in signal temporal logic. In *Proceedings of the 8th International Conference on Performance Evaluation Methodologies and Tools*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 66–73.
- [18] Luboš Brim, Karen Yorav, and Jitka Židková. 2005. Assumption-based distribution of CTL model checking. *International Journal on Software Tools for Technology Transfer* 7, 1 (2005), 61–73.
- [19] Luca Cardelli and Luis Caires. 2001. A spatial logic for concurrency. In *TACS*, Vol. 1. 1–37.
- [20] Luca Cardelli, Philippa Gardner, and Giorgio Ghelli. 2002. A spatial logic for querying graphs. In *Automata, Lang. and Program*. Springer, 597–610.
- [21] Otakar Cerba, Raitis Berzins, Karel Charvat, and Tomas Mildorf. 2016. Smart POI: Open and linked spatial data. In *EGU General Assembly Conference Abstracts*, Vol. 18. 12272.
- [22] Jingyu Chen, Ping Chen, Qiuyan Huo, and Xuezhou Xu. 2011. Clustering network-constrained uncertain trajectories. In *Fuzzy Systems and Knowledge Discovery (FSKD), 2011 Eighth International Conference on*, Vol. 3. IEEE, 1657–1662.
- [23] Qingping Chi, Hairong Yan, Chuan Zhang, Zhibo Pang, and Li Da Xu. 2014. A reconfigurable smart sensor interface for industrial WSN in IoT environment. *IEEE transactions on industrial informatics* 10, 2 (2014), 1417–1425.
- [24] Vincenzo Ciancia, Stephen Gilmore, Diego Latella, Michele Loreti, and Mieke Massink. 2014. Data verification for collective adaptive systems: spatial model-checking of vehicle location data. In *Self-Adaptive and Self-Organizing Systems Workshops (SASOW), 2014 IEEE Eighth International Conference on*. IEEE, 32–37.
- [25] Vincenzo Ciancia, Gianluca Grilletti, Diego Latella, Michele Loreti, and Mieke Massink. 2015. An experimental spatio-temporal model checker. In *International Conference on Software Engineering and Formal Methods*. Springer, 297–311.
- [26] Vincenzo Ciancia, Diego Latella, Michele Loreti, and Mieke Massink. 2014. Specifying and verifying properties of space. In *Theoretical Computer Science*. Springer, 222–235.
- [27] Vincenzo Ciancia, Diego Latella, Mieke Massink, and Rytis Pakauskas. 2015. Exploring spatio-temporal properties of bike-sharing systems. In *Self-Adaptive and Self-Organizing Systems Workshops (SASOW), 2015 IEEE International Conference on*. IEEE, 74–79.
- [28] Vincenzo Ciancia, Diego Latella, Mieke Massink, Rytis Paškauskas, and Andrea Vandin. 2016. A tool-chain for statistical spatio-temporal model checking of bike sharing systems. In *International Symposium on Leveraging Applications of Formal Methods*. Springer, 657–673.
- [29] Edmund M Clarke, Orna Grumberg, and Doron A Peled. 1999. *Model Checking*. MIT press.
- [30] Max J Egenhofer, Andrew U Frank, and Jeffrey P Jackson. 1989. *A topological data model for spatial databases*. Springer.
- [31] Max J Egenhofer and John Herring. 1990. Categorizing binary topological relations between regions, lines, and points in geographic databases. *The* 9 (1990), 94–1.
- [32] Sami Evangelista, Laure Petrucci, and Samir Youcef. 2011. Parallel nested depth-first searches for LTL model checking. In *International Symposium on Automated Technology for Verification and Analysis*. Springer, 381–396.
- [33] Wes Felter, Alexandre Ferreira, Ram Rajamony, and Juan Rubio. 2015. An updated performance comparison of virtual machines and linux containers. In *Performance Analysis of Systems and Software (ISPASS), 2015 IEEE International Symposium On*. IEEE, 171–172.
- [34] Martin Fowler. 2015. Microservices. (2015). Retrieved from: <http://martinfowler.com/bliki/MonolithFirst.html>.
- [35] Antony Galton. 2003. A generalized topological view of motion in discrete space. *Theoretical Computer Science* 305, 1 (2003), 111–134.
- [36] Martin Garriga. 2017. Towards a Taxonomy of Microservices Architectures. In *International Conference on Software Engineering and Formal Methods*. Springer, 203–218.

- [37] Alex Glikson, Stefan Nastic, and Schahram Dustdar. 2017. Deviceless edge computing: extending serverless computing to the edge of the network. In *Proceedings of the 10th ACM International Systems and Storage Conference*. ACM, 28.
- [38] S. Hendrickson, S. Sturdevant, T. Harter, V. Venkataramani, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau. 2016. Serverless Computation with openLambda. In *Proc. of the 8th USENIX Conf. on Hot Topics in Cloud Computing*. 33–39.
- [39] Mike Jia, Weifa Liang, and Zichuan Xu. 2017. QoS-aware task offloading in distributed cloudlets with virtual network function services. In *Proceedings of the 20th ACM International Conference on Modelling, Analysis and Simulation of Wireless and Mobile Systems*. ACM, 109–116.
- [40] Jessie Kennedy, Peter Barclay, and others. 1996. A survey of query languages for geographic information systems. (1996).
- [41] Jeffrey O Kephart and David M Chess. 2003. The vision of autonomic computing. *Computer* 36, 1 (2003), 41–50.
- [42] Edward A. Lee. 2008. *Cyber Physical Systems: Design Challenges*. Technical Report. EECS Department, University of California, Berkeley.
- [43] P. Leitner and J. Cito. 2016. Patterns in the Chaos – a Study of Performance Variation and Predictability in Public IaaS Clouds. *ACM Transactions on Internet Technology (TOIT)* 16, 3 (2016), 15.
- [44] Flavio Lerda and Riccardo Sisto. 1999. Distributed-memory model checking with SPIN. In *International SPIN Workshop on Model Checking of Software*. Springer, 22–39.
- [45] James Lewis and Martin Fowler. 2014. Microservices: A definition for this new architectural term. (2014). Retrieved from: <http://martinfowler.com/articles/microservices.html>.
- [46] Shancang Li, Li Da Xu, and Shanshan Zhao. 2015. The internet of things: a survey. *Information Systems Frontiers* 17, 2 (2015), 243–259.
- [47] Laura Nenzi, Luca Bortolussi, Vincenzo Ciancia, Michele Loreti, and Mieke Massink. 2015. Qualitative and quantitative monitoring of spatio-temporal properties. In *Runtime Verification*. Springer, 21–37.
- [48] Michael Roberts. 2016. Serverless architectures: What is serverless? (2016). Retrieved from: <http://martinfowler.com/articles/serverless.html>.
- [49] Y. Song, S. S. Yau, R. Yu, X. Zhang, and G. Xue. 2017. An Approach to QoS-based Task Distribution in Edge Computing Networks for IoT Applications. In *2017 IEEE International Conference on Edge Computing (EDGE)*. 32–39.
- [50] Josef Spillner, Cristian Mateos, and David A Monge. 2017. Faaster, better, cheaper: The prospect of serverless scientific computing and hpc. In *Latin American High Performance Computing Conference*. Springer, 154–168.
- [51] Haiying Sun, Jing Liu, Xiaohong Chen, and Dehui Du. 2015. Specifying Cyber Physical System Safety Properties with Metric Temporal Spatial Logic. In *2015 Asia-Pacific Software Engineering Conf*. IEEE, 254–260.
- [52] Christos Tsiganos, Timo Kehrer, and Carlo Ghezzi. 2016. Architecting dynamic cyber-physical spaces. *Computing* 98, 10 (2016), 1011–1040.
- [53] Christos Tsiganos, Timo Kehrer, and Carlo Ghezzi. 2017. Modeling and verification of evolving cyber-physical spaces. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2017, 2017*. 38–48.
- [54] Christos Tsiganos, Liliana Pasquale, Carlo Ghezzi, and Bashar Nuseibeh. 2018. On the Interplay Between Cyber and Physical Spaces for Adaptive Security. *IEEE Trans. Dependable Sec. Comput.* 15, 3 (2018), 466–480.
- [55] Christos Tsiganos, Liliana Pasquale, Claudio Menghi, Carlo Ghezzi, and Bashar Nuseibeh. 2014. Engineering Topology Aware Adaptive Security: Preventing Requirements Violations at Runtime. In *Proc. of the 22nd Int. Requirements Engineering Conf*. 203–212.
- [56] Monica Wachowicz, Rebecca Ong, and Chiara Renso. 2013. Tailoring trajectories and their moving patterns to contexts. In *Geographic Information Science at the Heart of Europe*. Springer, 285–303.
- [57] Jie Xu, Lixing Chen, and Pan Zhou. 2018. Joint Service Caching and Task Offloading for Mobile Edge Computing in Dense Networks. *arXiv preprint arXiv:1801.05868* (2018).
- [58] Zhixian Yan. 2010. Traj-ARIMA: A spatial-time series model for network-constrained trajectory. In *Proceedings of the Third International Workshop on Computational Transportation Science*. ACM, 11–16.
- [59] Zhixian Yan, Dipanjan Chakraborty, Christine Parent, Stefano Spaccapietra, and Karl Aberer. 2013. Semantic trajectories: Mobility data computation and annotation. *ACM Transactions on Intelligent Systems and Technology (TIST)* 4, 3 (2013), 49.
- [60] Ruozhou Yu, Guoliang Xue, and Xiang Zhang. 2018. Application Provisioning in Fog Computing-enabled Internet-of-Things: A Network Perspective. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 783–791.
- [61] Jing Yuan, Yu Zheng, Xing Xie, and Guangzhong Sun. 2011. Driving with knowledge from the physical world. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 316–324.
- [62] Jing Yuan, Yu Zheng, Chengyang Zhang, Wenlei Xie, Xing Xie, Guangzhong Sun, and Yan Huang. 2010. T-drive: driving directions based on taxi trajectories. In *Proceedings of the 18th SIGSPATIAL International conference on advances in geographic information systems*. ACM, 99–108.
- [63] Yu Zheng and Xiaofang Zhou. 2011. *Computing with spatial trajectories*. Springer Science & Business Media.