International Journal of Engineering Sciences Paradigms and Researches (IJESPR) (Vol. 47, Issue 04) and (Publishing Date: 10th December, 2018) (An Indexed, Referred and Impact Factor Journal) ISSN: 2319-6564 www.ijesonline.com A Self-stabilizing Control Plane for the Edge and Fog Ecosystems

Jnana Ranjan Sahoo¹, Swaha Pattnaik², Sulochana Nanda³, Suchismita Satpathy⁴

¹Department of Electronics & Communication Engineering, Raajdhani Engineering College, Bhubaneswar, Odisha
 ²Department of Electronics & Communication Engineering, NM Institute Engineering & Technology, Bhubaneswar, Odisha
 ³Department of Electronics & Communication Engineering, Capital Engineering College, Bhubaneswar, Odisha
 ⁴Department of Electronics & Communication Engineering, Aryan Institute of Engineering and Technology, Bhubaneswar, Odisha

Abstract

Fog Computing is now emerging as the dominating paradigm bridging the compute and connectivity gap between sensing devices (a.k.a. "things") and latency-sensitive services. However, as fog deployments scale by accumulating numerous devices interconnected over highly dynamic and volatile network fabrics, the need for self-configuration and self-healing in the presence of failures is more evident now thanever. Using the prevailing methodology of self-stabilization, we propose a fault-tolerant framework for distributed control planes that enables fog services to cope and recover from a very broad fault model. Specifically, our model considers network uncertainties, packet drops, node fail-stop failures and violations of the assumptions according to which the system was designed to operate, such as an arbitrary corruption of the system state. Our self-stabilizing algorithms guarantee automatic recovery within a constant number of communication rounds without the need for external (human) intervention. To showcase the framework's effectiveness, the correctness proof of the proposed self-stabilizing algorithmic process is accompanied by a comprehensive evaluation featuring an open and reproducible testbed utilizing real- world data from the intelligent transportation domain. Results show that our framework ensures a fog ecosystem recovery from faults in constant time, analytics are computed correctly, while the overhead to the system's control plane scales linearly towards the IoT load.

1 Introduction

Fog and Edge Computing are the technologies enabling computation at the network extremes, such as on downstream data, on behalf of cloud services, and upstream data, on behalf of IoT services [21]. The rationale of fog computing is that computing should happen at the proximity of the data source with the "fog" constituting any compute and network resources along the path between the data and the cloud. In this context, the "edge" differs from traditional sensing devices in that sensory data are processed in proximityand converted from raw signals to contextually relevant information [26].

In light of this, recent advancements in fog computing suggest using *cloudlets* as intermediate compute platforms between IoT devices (edge devices) and the cloud which allow users to exploit the analytic powerof the cloud without incurring the high latency in communicating with remote clouds [8]. A cloudlet (also referred as a foglet, gateway, microcloud) can be a single server or a small cluster of co-located servers that form a (virtual) pool of shared resources but from an external viewpoint are considered a single entity [17]. Compared to traditional datacenters, a cloudlet features much more limited resources, albeit its proximity to IoT devices makes it appealing for offloading compute tasks and receiving timely responses.

Although fog computing brings the computation closer to delay-sensitive services, the challenges re- stricting the cloud paradigm still remain as the pace of generated data continues to rise [25].

Now, these

overwhelming volumes of data not only have to be processed in time, but must be processed on, arguably, "weaker" hardware with potential nodes being vehicles, sensors, wifi access points, drones, cameras, and even wearable devices. Also, fog infrastructure usually operates in geo-distributed and less controlled environ- ments, with many applications competing for limited resources against high-priority services (*e.g.*, 5G) [24]. Consequently, failures due to hardware limitations and network uncertainties are highly likely at the fog continuum spanning between users, things, and clouds [5]. To maintain high availability, fog infrastructure must be resilient to both node and network failures. Thus, self-managing and self-healing solutions are required for fog ecosystems. IoT services must be able to recover from any issues that arise during their lifetime. In this context, it is critical to ensure continuous operation and recoverability at scale even in the event of failure without human involvement. In particular, cloudlets must satisfy the increasingly stringent fault-tolerance specifications of today's internet-enabled systems. In the current fog computing paradigm, fault-tolerance must be implemented to both preserve the system state locally at the edge and ensure the accuracy of analytics computations, especially in the case of a node failure or intermittent long-distance network connectivity problems.

We propose to address the challenge of dependable fog computing by using a fault-tolerant control planethat ensures service availability and data freshness in spite of the dynamic nature of the fog continuum. Via inter-connection of IoTs (edge devices), cloudlets and remote clouds, the proposed solution can tolerate network uncertainties, communication drops as well as cloudlet and IoT failures. In addition to these benign failures, our algorithms follow a very strong notion of fault-tolerance, called *self-stabilization* [13], which has provided the Internet with automatic failure recovery as early as the 1980's [19]. Self-stabilization ensures that the fog can recover after the occurrence of any temporary violations to the assumptions according towhich the system was designed to operate. These violations can include, for example, state corruption, extreme number of node failures, network partitions or unexpected system reconfiguration. Once such transient violations occur, non-self-stabilizing systems cannot guarantee correct system behavior due to data loss or the propagation of corrupted information. The correctness proof of a self-stabilizing system is required to guarantee recovery, within finite time, after the occurrence of the last transient violation.

Contribution and Research Outcome. This paper addresses the problem of how to tolerate and recover from run-time faults in distributed fog computing ecosystems. We consider a typical fog computing architecture, where edge devices are interconnected with remote clouds via network elements, denoted as cloudlets. Specifically:

- We introduce a self-stabilization framework for distributed control planes. The control plane is the core of the ecosystem and manages the network fabric with a global viewpoint and establishes the routing pathof data serviced by geo-distributed cloudlets. To the best of our knowledge, we are the first to introduce a self-stabilizing framework for control planes enabled over fog and edge ecosystems [18].
- To deal with a broad fault model that includes both communication and node failures, our correctness proof details how the proposed self-stabilizing solution can recover within a constant number of communicationrounds, after the occurrence of transient faults, as required by [12].
- To illustrate both the effectiveness and low runtime footprint of our framework at scale, we introduce a thorough evaluation using real-world data and actual queries of interest from an intelligent transportation service. Our results are reproducible and the reference implementation (including configuration and test data) is open-source and available online¹. Our experiments validate our analysis and show that even in the presence of severe failures, our solution can always recover in constant time while the network overhead scales linearly towards the IoT load.

Paper organization. Section 2 reviews related research. Section 3 presents the system model and objectives before proposing the solution for realizing the system in Section 4. Section 5 presents the correctness proof. Section 6 presents the experimentation, followed by the conclusion.



Figure 1: Table details the fault model and the chart illustrates when each fault set is relevant. The chart's gray shapes represent the system execution, and the white boxes specify the failures considered to be possible at different execution parts and recovery guarantees of the proposed self-stabilizing algorithm. The set of benign faults includes both packet failures and fail-stop failures.

2 Related Work

Fog and edge infrastructures are typically composed by hundreds of thousands to millions of heterogeneous and interacting components, which lead to the emergence of different types of faults. A major challenge in fog and edge computing is to define the fault and failure coverage required to provide high QoS [16]. Faults mayoccur either simultaneously or in any aspect of system operations ranging from application to hardware, and may have several causes, including insufficient memory, performance interference, system utilization, network congestion, server faults, application crashes, etc. Due to these challenges, existing work on fault- tolerance in large-scale distributed systems often have limitations in terms of practicality and performance guarantee. In [16], authors introduce CESSNA, a framework that provides consistency guarantees for stateful edge applications. CESSNA uses the Fault-Tolerant MiddleBox [20], which adopts the classical approach of "rollback recovery" where a system uses information logged during normal operation to correctly reconstruct state after a failure. In [27], authors present a fault-tolerant messaging architecture for edge systems. The fault-tolerance is achieved by introducing timing bounds that capture the relation between service parameters and loss-tolerance requirements. In [28], a fault-tolerant framework for data transmission in fog computing is introduced. The proposed fault-tolerance mechanism combines the advantages of Directed Diffusion andLimited Flooding to enhance the reliability of data transmission. We note that none of these solutions provides a holistic approach for addressing the fault-tolerance in edge and fog ecosystems.

Our framework fits naturally in distributed control planes, such as Istio and Linkerd [3, 4], that decouple operational control, policy enforcement and behavior telemetry from the business logic of distributed network fabrics and microservices. These frameworks provide fault-tolerance in the form of timeouts and (number of) retries for labelling nodes servicing HTTP requests as failed. In turn, circuit breaking is provided to safe-guard nodes overwhelmed by requests so that nodes "fail fast" when requests exceed the denoted limit. Thanks to our self-stabilizing algorithmic process, distributed control planes are introduced to a very strong notion of fault-tolerance on network uncertainties, communication drops, configuration errors, arbitrary transient violations, cloudlet and IoT fail-stop failures. In turn, no combination of faults can yield the system execution or corrupt data computations.

In the context of self-stabilizing algorithms and IoT, Siegemund *et al.* [22] present a self-stabilizing pub- lish/subscribe middleware for IoT applications. Their basic idea is that fault-tolerance is ensured through the construction of a distributed self-stabilizing data structure based on a virtual ring. However, operations over this ring take (*n*) time even in the absence of failures, where *n* is the ring size. Canini *et al.* [10] present self-stabilizing distributed control plane for software-defined networks (SDNs). Their work assumes that all nodes are either client hosts, switches or controllers. The algorithm stabilizes within $O(d^2n)$, where *d* is the

www.ijesonline.com

network diameter and *n* is the number of nodes. Chattopadhyay *et al.* [11] integrate an SDN control plane with the in-network processing infrastructure that can offload IoT services. They use a single centralized ser-vice deployment controller and lightweight SDN micro-controllers (μC). They mention that their algorithm for μC placement is self-stabilized with a linear convergence time (but no formal proof is provided). We provide both analytical and empirical proof for convergence in constant time. The state-machine replication technique used in this paper is inspired by practically-self-stabilizing virtual synchrony [14]. However, the proposed self-stabilizing solution has a much easier to understand leader election mechanism than the one in [14]. Moreover, our self-stabilizing solution stabilizes in constant time whereas the one in [14] does nothave a bounded stabilization time (by the definition of the solution criteria of practically-self-stabilizing systems).

While interesting and relevant, the above works do not address the impact of strong faulttolerance in ahierarchical network organization that includes cloud infrastructure, cloudlets that are placed at the networkedge and IoT devices. Our recovery time is within (1) and our placement mechanism convergence is within(1). We base our proofs on the definition of self-stabilizing systems [P13]. The definition requires the entire system to use bounded memory and recover after the occurrence of any transient violation of the assumptionsaccording to which the system was design to operate. To the best of our knowledge, this is the first work

that introduces a self-stabilizing control plane for the edge and fog ecosystems.

3 The System

Informatics is a science of abstractions, and a main difficulty consists in providing users with a "desired level of abstraction and generality — one that is broad enough to encompass interesting new situations, yet specific enough to address the crucial issues" [15]. This work provides a model that has the right-level of abstraction for the case edge computing since it allows both analytical and experimental study of the problem. We consider a fog computing system comprised of sets of nodes, such as the one of cloudlets C and IoT devices S, as well as a remote cloud infrastructure, which we refer to as the Cloud. Each cloudlet features specified communication, computation, and storage capabilities. Each cloudlet is associated with a wirelessaccess point covering a local area, referred to as a cell. The cloudlets in C form a shared resource pool that can serve the system collaboratively, e.g., aggregating IoT data and forwarding it to the Cloud. We assume that the cloudlets can share (over the Internet) such aggregated data with the Cloud by accessing a shared repository. The Cloud can use the repository to instruct cloudlets, *e.g.*, which queries the IoTs need to serve (edge devices), or provide advice the cloudlets on how to organize themselves, *e.g.*, propose the most-suitable leader according to the cloudlet specified capabilities and statistics gathered by the Cloud. The cloudlets themselves are intra-connected by backhaul links. We assume that, in the absence of failures, the quality of service of these links allow to send data and control messages in a timely manner — this is in contrast to the communications between the cloudlets and the Cloud, which we assume to be asynchronous by nature. The control plane manages and configures the cloudlets to route traffic and enforce service placement withIoT devices.

Objectives. We aim at developing a fault-tolerant framework for distributed control planes that enables large-scale fog services to cope well with communication uncertainties and a broad fault model without service downtime or the need for external (human) intervention. Next, we discuss the development objectives of the proposed self-stabilizing solution before specifying the system requirements.

- O1. The Cloud, cloudlets and IoTs (edge devices) should be able to exchange messages within a constantnumber of messages and communication rounds per information update.
- O2. The memory space and compute time of any system entity must always be bounded and networktraffic scale linearly to the number of system entities.
- O3. The presence of a constant number of benign faults (Figure 1) must not degrade the system perfor-mance beyond the bounds that are imposed by the system communication and processing delays. *I.e.*,



Figure 2: System overview



objective O1 must not be violated in the presence of benign faults (and the absence of violations consideredin objective O4).

- O4. We also consider arbitrary transient violations of the assumptions according to which the system wasdesigned to operate (as long as the algorithm code stays intact). After the occurrence of these violations, the system must recover autonomously within a constant number of communication rounds and return to satisfy the task specifications. By autonomous we mean the absence of external intervention (of a humanor a system component that is not part of the proposed framework).

Specifications. The control plane for the edge organizes the *cloudlet layer* (Figure 2), such that in the presence of communication and node failures cannot disrupt the execution of services, such as IoT queries. In detail, we require the implementation of the following functionality:

i) The cloudlet and IoT registration allows the Cloud to include individual nodes in the system (Figure 2). A node is allowed, after a predefined delay and local cleanups, to register again when it notices that it became disconnected from the system due to failures. Note that the latter case is rare, and thus, it shouldnot repeatedly consume significant system resources.

(ii) The query functionality allows the Cloud to request the flow of information according to a model that the IoTs (edge devices) are to update periodically. That is, given the Cloud's current belief about the query result, the specified IoTs (edge devices) will update the system whenever the collected sensory information deviates from the model. The cloudlet aim here is to aggregate these updates so that a concise query result arrives to the Cloud. Since this needs to be done in the presence of communication and node failures, each IoT should send its updates to a set of cloudlets and the latter should acknowledge (Figure 3). The cloudlets then should use a *leader* to unify their updates and forward concise query results to the Cloud. The cloudlet layer must function well in case of a failing leader. Therefore, a set of cloudlets, called guards, should monitor he leader's activity and guarantee query result delivery until the system decides on a new leader (Figure 4).

(iii) The management of general purpose services can help to overcome capability differences among individual nodes via task load-balancing. Such tasks can be initiated by IoT users that need to leverage

International Journal of Engineering Sciences Paradigms and Researches (IJESPR) (Vol. 47, Issue 04) and (Publishing Date: 10th December, 2018)

(An Indexed, Referred and Impact Factor Journal)

_______ISSN: 2319-6564



Figure 4: The inner-structure of the proposed cloudlet layer

Algorithm 1: A high-level overview on algorithms 3 and 4 1 Registers shared between the modules in algorithms 3, 5, and 6 info: has the form of (devices, cloudlets, leader, guards), where the field *devices* is a set of IoT devices, their models and the information needed for failure detection; cloudlets is a set of cloudlets and the information needed for failure detection; leader of the form (seq, id) is the cloudlets' current leader and an associated sequence number; guards is a set of cloudlets ids (a subset of cloudlets) that have been selected as guards; */ the module for the self-stabilizing cloud (Algorithm 3) 2 Local variables: new Cloudlet/new Iot: new cloudlets and IoTs (edge devices) and their models; sequence: leadership number; 3 do forever begin if the reset procedure is inactive and fresh information was recived from all trusted (not to be faulty) cloudlets 4 then Use newlot, newCloudlet and fault detection information for updating devices and cloudlets, respectively; 5 6 if leader / cloudlets then elect a leader with sequence++; if *sequence* = *MAXINT* then invoke the reset procedure; 7 if guards \cap cloudlets = \emptyset then select new guards; 8 Once the reset procedure is done, initialize all local variables; Q 10 upon registration request arrival from an IoT or a cloudlet, update the set *newIot* and *newCloudlet*, respectively; 11 upon RESET message arrival, invoke the distributed reset procedure; the module for IoT (Algorithm 4) 12 Local variables for the IoT module: model: a data structure that encodes the recent sensory readings; cloudletModel: recent model received from the cloudlet; cloudletList: dissemination point list; lastUpdate: time of the last update reception from a cloudlet; *msgseq*: a positive integer used as a sequence number for messages sent tocloudlets; *MSG*: a set that stores the highest message sequence received; 13 do forever begin if lastUpdate was too long ago then initialize and register at the Cloud this IoT; 14 else if an update is needed then 15 foreach cloudlet in cloudletList do send msgs &q++, model; > 16 if *msgseq* = *MAXINT* then invoke the reset procedure; 17 18 upon $m = \langle seg, list, model \rangle$ arrival from a cloudlet do { if m's seg is fresh then update cloudletList, cloudletModel, *lastUpdate* and *MSG* reply;} **19 upon** reply m = (seq) arrival from a cloudlet, update *msgseq*;

on the cloudlet capabilities. Also, cloud services may wish to avoid communication-intensive computations, such as virtual traffic light that base its decisions on the current road traffic conditions that different vehicles report. The fault-tolerant management of such services can be based on state-machine replication that is well-synchronized with query operations.

4 Proposed Solution

Algorithms 1 and 2 provide a high-level description of our solution, and the details appear in algorithms 3, 4, 5 and 6, which implement the proposed solution to the above task specifications by considering the code to be executed by the Cloud, IoT devices, cloudlets, and respectively, the emulators of the replicated state-machine.

<u> </u>	anithme 2. A high land an anithma 5 and 6
Alg	orithm 2: A high-level overview on algorithms 5 and 6
/*	the module for the self-stabilizing cloudlet (Algorithm 5) */
20 Loc S ['] tl (i tl	al variables: <i>deviceSet</i> : a set of IoT devices and their most recently received models; <i>agreegateInfo</i> : a set of data tructures encoding aggregated sensory information; <i>msgc</i> : a positive integer used for ordering message sent to he leader and guards; <i>msgtoiot</i> : a positive integer used for ordering messages sent to IoT devices; <i>MSGc</i> : a set of <i>id</i> , <i>seq</i>) pairs that stores the highest message sequence received by cloudlet <i>id</i> ; <i>MSGSEQ</i> : a set of <i>(id, seq)</i> pairs that stores the highest message sequence received by IoT <i>id</i> ;
21 do 1	forever begin
22 23	if the reset procedure is inactive and i / cloudlets then initialize and register at the Cloud this cloudlet; else if the reset procedure is inactive then
24 25	Use <i>devices, deviceSet</i> and <i>cloudlets</i> to update <i>deviceSet, MSGSEQ</i> and <i>MSGc</i> , respectively; foreach <i>IoT j that this cloudlet is responsible for</i> do send info about <i>msgtoiot++</i> , the cloudlets that are responsible for this IoT and the model of this IoT;
26	foreach <i>j</i> that is a guard or a leader do send info about msgc++ and the aggregated data received from the lo'T that this cloudlet is responsible for;
27	If $msgseq = MAXINI$ then invoke the reset procedure;
28 upo 29 upo 30 upo	on $\overline{m} = \langle seq, model \rangle$ arrival from an IoT do {Acknowledge <i>m</i> and update <i>deviceSet</i> and <i>deviceSet</i> } on $m = \langle seq, aggregated \rangle$ arrival from a cloudlet do {Acknowledge <i>m</i> and update <i>agreegateInfo</i> and <i>MSGc</i> } on $m = \langle seq \rangle$ arrival from an IoT do {update <i>msgtoiot</i> } on $m = \langle seq \rangle$ arrival from a cloudlet do {update <i>msgcoiot</i> }
/*	m = (seq) and a leader (Algorithm 6) (update msgc)
32 Loc n si u (1	variables: <i>replicaState</i> []: an array of the state machine's replica, where $rep[i]$ refers to the one that processorp _i maintains, and <i>replicaState</i> [j] refers to the last arriving message from p_j containing p_j 's <i>replicaState</i> [j]. <i>myLeader</i> tores the identifier of the local leader. The term view refers to the set of replicas that the leader considers to be p and connected, <i>i.e.</i> , they can participate in the emulation of the state-machine. <i>FD</i> stores the processors that the local) failure detector considers as active;
33 do	forever begin
34	if the Cloud propses leader to be this replica but myLeader does not or the view is not all trusted (not to be failing) guards and this replica then propose a view with this replica as a leader as well as all trusted (not to befailing) guards as members:
35 36	if myLeader refers to this replica but the Cloud proposes another trusted guard then update myLeader to the proposed one; if this replica is the leader and all replicas have completed a communication round then compute the new state of the automaton and update replicaState:
37 38	else update <i>replicaState</i> according to the one of the leader replica and send your input to the leader; if this replica is a guard that is not lLeader but lLeader is suspected (to be failing) then reset myLeader and update data about the local state of this replica:
39	else if the myLeader is well-defined (not reset) then send this replica's state to leader;
40	if this replica is the lLeader then broadcast this replica's state to $lGuards \cap FD$;
41 up	on <i>m</i> arrival from a guard or a leader do {update <i>replicaState</i> with <i>m</i> ;}
Algon Over devic cloud senso to the	withm 3 assumes the availability of a self-stabilizing cloud infrastructure, such as [7]. wiew. The Cloud periodically monitors the system and keeps track of the Cloudlets and IoT we sthat are up and running. Based on this information, and according to some mapping, each let is associated with a list of IoT devices. The IoT devices periodically send their data (e.g., bry information) to their associated cloudlet(s). Instead of having each cloudlet to report directly the cloud, each cloudlet reports its collected data to a leader. The leader is the one that collects

cloudlet is associated with a list of IoT devices. The IoT devices periodically send their data (e.g., sensory information)to their associated cloudlet(s). Instead of having each cloudlet to report directly to the cloud, each cloudlet reports its collected data to a leader. The leader is the one that collects and aggregates all data and reports it to the Cloud (via shared registers). The above constitutes a "normal"(fault-free) operation. However, due to unexpected transient faults or more permanent faults (e.g., a cloudlet fail-stopping), as well as the need for bounded counters, additional checks must take place at the different components of the system. Algorithms 3, 4, and 5 present such details for the Cloud, the IoT devices and the Cloudlets, respectively. Furthermore, in the event that the leader fail-stops, we do not want the data flow to the cloud to be suspended or critical information to be lost. To this respect, from the list of operational cloudlets, the Cloudalso appoints a set of guards. The purpose of the guards is to monitor more frequently the status of the leader and in the event that the leader fail-stops, they report the latest collected data to the Cloud. Therefore, in Algorithm 5, each cloudlet reports its collected data not only to the leader, but also to the guards. Since the leader and the guards need to maintain consistent information on the collected data (and on any other

ISSN: 2319-6564

www.ijesonline.com

information the control plane could be maintaining), they run Algorithm 6, which realizes a selfstabilizing state-machine replication mechanism. In Section 5 we provide the correctness proof illustrating that our algorithmic framework can self-stabilize in a constant number of communication rounds, while Section 6 shows through a large testbed that there is no information loss even in the presence of multiple, different and randomly injected failures to the fog ecosystem.

We now proceed to present more details. We start by describing the registers that are shared by the nodes. Then, we go through the code according the above functionality list.

Registers. The shared register *data* stores the aggregated sensory information that is collected by the IoTs (lines 65–66), aggregated by their corresponding cloudlets (line 97), and written by the leader (line 125). The Cloud and the cloudlet exchange control information via the shared registers *info* and *infoAck*. The register *info* includes the fields (IoT) *devices, cloudlets, leader* and *guards*. The register *infoAck* is an array, such that the entry *infoAck* [k] holds $p_{k's}$ acknowledgment, where $p_k \in C$ is a cloudlet and the acknowledgment includes all the fields of *info*. In detail, the Cloud, $p_{cloudID}$, stores its view on the system membership in*info* (line 56) and cloudlet p_k acknowledges the reception of this information by copying the value of *info* to *infoAck* [k] (line 88). Moreover, $p_{cloudID}$ selects, when needed, new cloudlets' leader (line 51) and guards (line 53).

Registration. IoTs (edge devices) and cloudlets register directly at the Cloud by sending a registration message (lines 65 and 89) after initializing their local variables and communication channels. This initialization guarantees that the joining node (or its communication channels) does not hold stale information. Once the registration message arrives to the Cloud, $p_{cloudID}$, the Cloud lists the joining node as a newcomer (lines 57 and 58). These newcomers will be listed as the system's IoT devices and cloudlets (lines 49 to 50) after the completion of the previous update round of these sets, which line 48 assures. The proposed solu-tion assumes access to unreliable failure detectors. This allows the Cloud not to wait for cloudlets that are suspected to be faulty as well as to remove failing nodes from the IoT and cloudlet sets.

Query. We consider queries that are initiated by Cloud applications and require repeated updates. These queries include the Cloud current belief about the anticipated result, which we refer to as the query model. This allows IoT devices to reduce the number of times in which they transmit results to periodic queries since there is no need to transmit a result that fit the current belief of the Cloud according to the query model.

In detail, the registration procedure constructs up-to-date views on the sets of IoT devices and *cloudlets* in the shared register together with the current leader and guards. The proposed solution associates with each IoT the query description and model. This information is stored in devices. The cloudlets use a function, myIoT(), for mapping between them and the IoTs that they are responsible to communicate with (line 96). (A possible mapping could be to have the IoTs being assigned to the cloudlets in the same region, basedon their proximity. Nevertheless, our system is independent on the specific mapping employed.) Cloudlets send the queries (along with their models) to these IoTs. The latter store the arriving information and acknowledge (lines 71 to 74). Once in a predefined periodicity, the IoTs update the query results, if needed (line 66). The cloudlets acknowledge the update arrival (lines 74 and 104). The cloudlets in turn periodically aggregate the sensory information received by the IoTs and send it to the leader and the guards (line 97). The leader updates the shared repository with the query results (line 125), whereas the guards serve as warm-backup leaders. We assume access to the functions *electLeader()* and *selectGuards()* that for a given set of system cloudlets elect a leader and select guards, respectively. In electing a leader and guards, we may want nodes that are more stealth, maybe closer to the IoT devices or in the center of the coverage area (e.g., in the center of the city); the leader/guard selection problem can be inherent to the fog service placement problem (FSPP) [23], which is a different challenge in fog computing than the studied one. Nevertheless, inour system we could swap in/out FSPP algorithms and we are resilient to the algorithm in use.

State-machine replication. Since both the leader and the guards receive aggregated sensory information from the cloudlets, they need to be in sync with respect to this information. More generally speaking, the leader and the guards could provide additional service as part of the control plane. So, they need to coordinate their activities and maintain consistent state between them. The fact that the system is asynchronous, together with the need for self-stabilization, makes it quite a challenging task. To this respect, we have the

www.ijesonline.com

leader and the guards to run Algorithm 6.

The algorithm maintains a consistent state (aggregated sensory information) by performing multicast rounds coordinated by the leader. All necessary replica information (including the state) is maintained by each node in array *rep*[] (line 113), which is exchanged between the leader and the guards (lines 143-145). In detail, once a cloudlet realizes that it has become the leader (line 130), it proposes to install a view of the current members, which includes itself and the guards that according to its local failure detector havenot fail-stopped. The guards start following the leader towards installing this view by adopting its proposal (line 140). Once the leader sees that the view members have adopted its proposal (lines 133 and 115), itbuilds the new state based on the collected messages and states (lines 136 and 116) and proceeds to install the view. The guards adopt the leader's *rep* – including the (new) state (lines 139 and 120) completing in this way the installation of the view (lines 135 and 117). The multicast rounds can now begin, which are coordinated by the leader (lines 134 and 122-126) and followed by the guards (lines 138 and 121). The access to the application's message queue (commands to be executed by the state machine) is done via fetch(), which returns the next multicast message; the state transition function *apply(state, msg)* applies the aggregated input array *msg* to the replica's *state* and produces the local side effects. Simply put, in our case, the input to the state machine is the aggregated sensory information, which is sent by the cloudlets to the leader and the guards in Algorithm 5 (line 97) and stored by the latter in agrregateinfo (line 107). So, essentially the multicast rounds of the state machine keep this information consistent among the leader and the guards. At the end of each multicast round, the leader updates the sensory information maintained in the shared register data (line 125).

In the event of a leader fail-stop, and until the Cloud assigns a new leader (line 51), the guards update the *data* repository (lines 141-141), instead; this ensures a continual update of the sensory information (which, depending on the application, could be crucial). If there is a change in the set of guards (either due to a fail-stop or due to an update of this set by the Cloud), then the leader begins the procedure to install a new view (line 131) with the new membership, without the need of any external intervention (including that of the Cloud). The failure detector abstraction (defined in line 113) can be implemented using heartbeats and counter thresholds (see for example [9]), or using "hello" messages and timeouts in a more time-informed setting (as we do in our simulation study in Section 6).

Recovering the system state via global reset. Self-stabilization requires bounded space, which includes bounded counters. Counters can grow up to a predefined size *MAXINT*, *e.g.*, 2^{64} 1. Under normal operation, and if say, a counter is incremented every nano-second, then this limit could be reachedin approx. 146 years. However, a transient violation of the assumptions according to which the system was designed to operate can corrupt the counter and cause it reach *MAXINT*. In such a case (lines 68, 99, and 144), the cloudlet or IoT holding this counter will send a RESET message to the Cloud, calling for a global system reset. The Cloud, upon receiving such a message (line 59) or the *sequence* counter reaches *MAXINT* (line 52), initiates the *reset procedure*: it sets the shared register *info* into (line 52 or line 59), and waits until all non-faulty cloudlets have acknowledged this (via the shared array *infoAck*, line 55), beforeit unregisters all cloudlets and IoT devices (by setting *info* into (, , ,)) and flashes all its local variables. This causes each cloudlet (line 89) to register again after a local reset of the node state and its communication channels, following the registration procedure@des@mbed above. Since the IoTs are no longer in *info.devices*, no cloudlet will contact them, causing each (non-faulty) IoT to timeout and hence also register again aftera similar initialization procedure (line 64).

5 Correctness Proof

Our analysis demonstrates a constant time recovery from arbitrary transient faults. It considers the *interleaving model* [13], in which the node's program is a sequence of *(atomic) steps*. Each step starts with an internal computation and finishes with a single communication operation, *i.e.*, message *send* or *receive*. The *state*, s_i , of node $p_i \in P$ includes all of p_i 's variables as well as the set of all incoming communication channels. Note that p_i 's step can change s_i as well as remove a message from *channel_{j,i}* (upon message arrival) or add a message in *channel_{i,j}* (when a message is sent). The term *system state* refers to a tuple of

Algorithm 3: Code for the self-stabilizing cloud *p*_{cloudID}.

- 42 Variables: new Cloudlet/new I of: new cloudlets and IoTs and their models (bounded by cloudletSetSize); sequence: leadership number:
- 43 Shared registers: *data*: is a data structure that stores the sensory information, to be processed by the cloud depending on the application; it includes records of the form (*id, leader, round, dat*), where *id* is the cloudlet's unique id that included the context *dat* in the data structure, at round *round* of the state machine with leader *leader*; *info*: has the form of (*devices*, *cloudlets*, *leader*, *guards*), where the field *devices* is a set (bounded by *deviceSetSize*) of IoT devices, their models and the information needed for failure detection; *cloudlets* is a set (bounded by *cloudletSetSize*) of cloudlets and the information needed for failure detection; *leader* of the form (seq, id) is the cloudlets' current leader and an associated sequence number; guards is a set of cloudlets ids (a subset of cloudlets) that have been selected as guards; infoAck[cloudletSetSize]: an array that stores the latest value of info that each cloudlet has read;
- 44 Interface: suspectedIot(set) and suspectedCloudlet(set): return the sets of suspected to be faulty IoT devices and cloudlets, respectively; *electLeader(set)*: returns the elected leader from *set*; *selectGuards(set)*: returns the set of guards from set

45 do forever /* use predefined periodicity */ begin

- let lInfo := (lDevices, lCloudlets, lLeader, lGuards) := read(info); 46
- 47 let lInfoAck := read(infoAck);
- if $\perp \neq lInfo$ (lanfo $\neq lIafoAck[k]: k \in C \suspectedCloudlet(C)$) then 48
- $(lDevices, newlot) \leftarrow ((lDevices \setminus \{(k, \bullet) : k \in suspectedIot(lDevices))\}$ newlot,); Ø 49
- $(|C|| coullets, newC|| coullet) \leftarrow ((|C|| coullets \setminus \{(k, \cdot) : k \in suspectedC|| coullet(|C|| coullets))) ewC|| coullet,);$ Ø 50
- 51 if *lLeader.id* $/ \in lCloudlets$ then *lLeader* \leftarrow (sequence++, electLeader(lCloudlets));
- if sequence = MAXINT then write(info,); \perp 52
- **if** (*lGuards Cloudlets*) = **then** *lGuards* ← selectGuards(lCloudlets $\ \{lLeader.id\};\$ 53
- write(info, (lDevices, lCloudlets, lLeader, lGuards)); 54

else if $\{\perp, (\emptyset, \emptyset, \emptyset, \emptyset)\} \supseteq (\{ llnfo \} \cup \{ llnfoAck[k] : k \in C \setminus suspectedCloudlet(C) \})$ then write(info, $(\emptyset, \emptyset, \emptyset, \emptyset)$); (newCloudlet, newIot, sequence) $\leftarrow (\emptyset, \emptyset, 0)$; 55

- else if $\{\bot\} \subset (\{lInfo\} \cup \{lInfoAck[k] : k \in C \setminus suspectedCloudlet(C)\})$ then write(info, \bot); 56
- 57 upon message $m = \langle \text{REGISTER} \rangle$ arrival from IoT *j* at time *t* do newlot \leftarrow (newlot $\cup \{(j, t, \bot)\}$);

ss upon message $m = \langle \text{REGISTER} \rangle$ arrival from cloudlet z at time t do newCloudlet \leftarrow (newCloudlet $\cup \{(z, t)\}$);

<u>59 upon message $m = \langle \text{RESET} \rangle$ arrival from device k do write(*info*, \perp);</u>

the form $c = (s_1, s_2, \dots, s_n)$ (system configuration), where each s_i is p_i 's state (including messages in transit to p_i). An execution (or run) $R = c_0, a_0, c_1, a_1, \dots$ is an alternating sequence of system states c_x and steps a_x , such that each c_{x+1} , except c_0 , is obtained from the preceding one, c_x , by the execution of step a_x . We say that execution R is legal if it satisfies the task specifications throughout R. We say that a system state c is safe if every execution that start from c is legal. Definition 5.1 considers a system state that Theorem 5.1 shows to be safe.

Definition 5.1 (Safe system state) We say that the system state c is safe if the following hold. (1) Let $p_i \in C$ and $p_i \in S$, such that $(j, m_j) \in myIoT$ (devices_{cloudID}, cloudlets_{cloudID}). It holds that cloudletList_j = cloudletList(j,

 $cloudlets_{cloudID}$) \wedge ($lastUpdate_j \leq clock_j$ ()). Moreover, $devices_{cloudID} = \{(k, \bullet) \in deviceSet_i\} \wedge ((z, t, \bullet) \in agreegateInfo_j)$ $\Rightarrow p_z \in C \land t \leq clock_i()) \land (j, \bullet, m_i) \in agreegateInfo_i$ (2) The value of msgseq_i msgc_i and msgtoiot_i is greater or equal to any value of msgseq, msgc, and respectively, msgtoiotfields associated with p_i in messages and cloudlets. (3) |A| = 1, where $A = \{(v, su, r, sa, m) : p_i, p_j \in$ $C \land (v, su, r, sa, m, \bullet) = r \in \{rep_i[j], rep_{i,j}\}\}$, such that $rep_{i,j}$ is a message that was sent in line 143 from p_i

to p_j . Moreover, $msg_{leader_{cloudID}.id}[k] = input_k$, where $p_k \in leader_{cloudID} \cup guards_{cloudID}$. (4) No counter has reached MAXINT and there are no $\langle RESET \rangle$ messages.

We say that an execution is *fair* if every step that is applicable infinitely often is executed infinitely often. Theorem 5.1 demonstrates the required properties for self-stabilization and use the term (asynchronous) cycles of a fair execution R. A cycle is the shortest prefix of R in which every non-failing node p_i performs a completed iteration of node p_i 's do forever loop, all messages that p_i sent during that iteration were delivered, and all of the iteration's requests were replied.

Theorem 5.1 *The system's state is safe within* O(1) *cycles.*

60 Local state: model: a data structure that encodes the recent sensory readings; cloudletModel: recent model received from the cloudlet; cloudletList: a list (bounded by cloudletListSize) of dissemination points (ordered by descending priority); lastUpdate: time of the last update reception from a cloudlet (according to IoT's local time); msgseq: a positive integer used as a sequence number for messages sent to cloudlets; MSG: a set of (id, seq) pairs that stores the highest message sequence received by cloudlet id;
61 Interface: update(): receives the last sent model and received cloudletModel as well as the time in which that reception occurred (lastUpdate). The function then updates model (and returns true) if the cloudlet model requires an update due to change in sensory input, a timeout due to a missing acknowledgment from the cloudlet or a changein the cloudlet model specifications;
62 Function: <i>iotInit</i> (): the IoT device first resets all variables dealing with Cloudlet data and control information as well as local data and control variables. Then it sends a special message <i>INIT</i> to the Cloud, so that the Cloud removes all information about this device from the Cloudlets. Once this is done, the Cloud returns an acknowledgment to the device, and the function returns.
63 do forever /* use predefined periodicity */ begin
if (clock() —lastUpdate) > LIMIT then IoT init(); send(cloudID, REGIS(TER); >
65 else if update(model, cloudletModel, lastUpdate) then
66 for each $id \in cloudletList$ do send(id , (msgseq, model));
67 $msgseq \leftarrow msgseq + 1 /* \text{ if a message was sent }*/$
68 _ if $msgseq = MAXINT$ then send($cloudID$, (RESET));
⁶⁹ upon $m = (seq, list, model)$ arrival from cloudlet j at time $t = clock()$ begin
70 if $m.seq > MSG j.seq$ then
71 $((cloudletList, cloudletModel), lastUpdate) \leftarrow ((m.list, m.model), t);$
$nsg \leftarrow (MSG \leftarrow (MSG \setminus \{(k, \bullet) : k \in cloudletList \lor k = j\}) \cup (j, m.seq);$
rs $send(j, \langle MSG j.seq \rangle);$
74 upon message $m = \langle seq \rangle$ arrival from cloudlet z do msgseq $\leftarrow \max\{m.seq, msgseq\}$;

Algorithm 4: Code for IoT iot

Proof. The proof considers the predicate $pred = \{\bot\} /= (\{info\} \cup \{infoAck[k] : k \in C \setminus suspectedCloudlet_{cloudID}(C)\})$. We start by considering an execution in which *pred* holds throughout *R*, and thus $p_{cloudID}$ does not executes lines 55 and 56. Under this assumption, we show that items 1 to 4 of Definition 5.1 hold within (1) cycles. As a completely case, we consider a starting system state in which *pred* does not hold, and show that, within O(1) cycles, the *pred* holds.

Item 1. Let $(p_i, p_j) \in C \times S$. Within O(1) cycles, the cloud $p_{cloudID}$ updates the *devices* and *cloudlets* fields in *info* (line 54). Within O(1) cycles, p_i reads *devices* and *cloudlets* (line 88) and send $\langle \bullet, cloudletList(j, lCloudlets), m \rangle$ to IoT p_j (line 96), such that $(j, m_j) \in myIoT(devices_{cloudID}, cloudlets_{cloudID})$. When that message arrives, p_j stores it in *cloudletList_j* and *cloudletModel_j* as well as updates *lastUpdate_j* with the arrival time (line 71). Thus, *cloudletList_i* = *cloudletList(i, cloudlets_{cloudID})* \land (*lastUpdate_i* \leq *clock_i*()). Lines 92 and 102 implies $devices_{cloudID} = \{(k, \bullet) \in deviceSet_j\}$ and line 107 implies $((z, t, \bullet) \in agreegateInfo_j \implies p_z \in C \land t \leq clock_i()) \land (j, \bullet, m_j) \in agreegateInfo_j$.

Item 2. Suppose that in R'_{i} starting state, Item 3 does not hold with respect to a $p_i'_{i}$ field. Within O(1) cycles, any message containing *msgseq*, *msgtoiot* or *msgc* arrive to its destination p_j . Thus, for the sake of a simple presentation, we focus on the case in which Item 3 does not hold in node p_j with respect to a field that is associated with node p_i . We observe that within (1) cycles, p_i and p_j complete a message round-trip that include this filed. In detail, these message are sent in lines 66, 96 and 97 and received inlines 74, 110 and 111, respectively. Note that whenever p_i receives any such message, p_i updates the local value with the received one, in case the latter is greater than the former.

Item 3. Within (1) cycles, *leader* and *guards* are set by $p_{cloudID}$ (line 54) and all cloudlets read these values (line 88). We show that if a new leader has been put in place (or the view has become inconsistent), the leader installs a new view which includes itself and the guards. For this purpose, it first proposes thisview (line 136), which it is accepted by the guards (line 140) within (1) cycles. Then, within (1) cycles it updates rep[] and installs this view (lines 135 and 139), in which the guards have updated their rep[]s based on the one of the leader. After that the leader resumes the round-base updates for maintaining the state among itself and the guards (lines 134, 138, and 122–126), hence stabilizing the state machine replication. Moreover, it aggregates $input_k : p_k \in leader_{cloudID} \cup guards_{cloudID}$, such as $msg_{leader_{cloudID},id[k] = input_k$.

Algorithm 5: Code for cloudlet p_i
 75 Local state: deviceSet: a set (bounded by deviceSetSize) of IoT devices and their most recently received models; 76 agreegateInfo: a set of data structures encoding aggregated sensory information; 77 msgc: a positive integer used for ordering message sent to the leader and guards; 78 msgtoiot: a positive integer used for ordering messages sent to IoT devices; 79 MSGc: a set of (id, seq) pairs that stores the highest message sequence received by cloudlet id; 80 MSGSEQ: a set of (id, seq) pairs that stores the highest message sequence received by IoT id; 81 Shared registers: info and infoAck: as in Algorithm 3; 82 Interface: aggregate(deviceSet): returns the aggregated sensory information:
 ⁸³ cloudletList(k, set): for a given IoT device <i>iot_k</i> and a <i>set</i> of cloudlets, this function returns the cloudlet list that <i>iot_k</i> should use (prioritized in an descending order); ⁸⁴ myIoT(): projection of the IoTs that are within the cloudlet's responsibility; ⁸⁵ cloudID: the address of the Cloud;
<i>so function: cloudermit():</i> the clouder instresses all variables dealing with the data and control information of cloudlets and IoT devices as well as its local data and control variables. Then it broadcasts a special message <i>INIT</i> to all other cloudlets, and to the Cloud so that the other cloudlets remove all information about this cloudlet; the Cloud removes all relevant information about this cloudlet from the IoT devices. Once the cloudlet receives acknowledgments from all the cloudlets and the Cloud, the function returns.
87 do forever /* use predefined periodicity */ begin
ss let $lInfo := (lDevices, lCloudlets, lLeader, lGuards) := read(info); write(infoAck[i], lInfo);$
89 II $lingo \neq \pm N/\Theta$ loudiets then $cloquaterinit()$; send(cloualD, REGIS(ER)); } 90 else if $linfo = l + then$
if $i/\notin Guards \cup \{ILeader.id\}$ then $(agreegateInfo, MSGc) \leftarrow \emptyset, \emptyset$:
92 $deviceSet \leftarrow (deviceSet \setminus \{(k, \cdot): k / \notin Devices); \}$
$MSGSEQ \leftarrow (MSGSEQ \setminus \{(k, \cdot): k / deviceSet); \}$
94 $MSGc \leftarrow (MSGc \setminus \{(k, \cdot): k \mid \notin cloudlets); \}$
95 let (iotAdd, msgAdd) := (0, 0); 96 foreach (j, m) \in myIoT(lDevices, lCloudlets) do {send(j, {msgtoiot, cloudletList(j, lCloudlets), m}); 107 iotAdd \leftarrow 1};
foreach $j \in lGuards \cup \{lLeader.id\}$ do $\{send(j, (msgc, aggregate())); msgAdd \leftarrow 1\};$
98 $(msgtoiot, msgc) \leftarrow (msgtoiot + iotAdd, msgc + msgAdd);$
if $MAXINT \in \{msgc, msgtoiot\}$ then send(cloudID, (RESET));
100 upon message $m = \langle seq, model \rangle$ arrival from IoT j at time t begin
101 if $m.seq > MSGSEQ j.seq$ then
$102 \qquad deviceSet \leftarrow (deviceSet \setminus \{(j, \bullet)\}) \cup \{(j, t, m)\};$
$MSGSEQ \leftarrow (MSGSEQ \setminus \{(j, \bullet)\}) \cup (j, m.seq);$
104 $\operatorname{send}(j, \langle MSGSEQ j.seq \rangle);$
105 upon message $m = (seq, aggregated)$ arrival from cloudlet z at time t begin
106 if $i \in Guards \cup \{ Leader.id\} \land m.seq > MSGc z.seq$ then
$107 \qquad agreegateInfo \leftarrow (agreegateInfo \setminus \{(z, \bullet)\}) \cup \{(z, t, m)\};$
$\frac{108}{108} = MSGC \leftarrow (MSGC \setminus \{(z, \bullet)\}) \cup (z, m.seq);$
109 Send(z , (MSGc(z , seq));
110 upon message $m = (seq)$ arrival from 101 k do msgrotor $\leftarrow \max\{m.seq, msgrotor\}$
In upon message $m = \{seq\}$ arrival from clouder ζ up $msge \leftarrow \max\{msge\}$
Item 4. Suppose that in R_i starting state, Item 3 does not hold at node p_i . We observe that within $O(1)$ cycles, either Item 3 holds and Item 4 does not hold, or both hold. Moreover, within $O(1)$

additional cycles, RESET arrives to $p_{cloudID}$ and the assumption above does not hold (line 59). Thus, the rest of the proof considers complementary case.

For the case that is competently to the assumption that appears in the proof start, suppose, For the case that is competently to the assumption that appears in the proof start, suppose, towardsa, contradiction, that pred does not hold in R's starting state. Moreover, suppose that any prefix R of R = R R that has (1) cycles, does not have a matching suffix R during which the predicate pred holds. Since pred does not hold during R, $p_{cloudID}$ does not executes lines 49 to 54 during R. Therefore, it must executes either line 55 or 56 for a constant number of times during R. Suppose that $p_{cloudID}$ does not execute line 55 during R. Thus, within (1) cycles, $p_{cloudID}$ executes repeatedly line 56 until the if-statement condition of line 55 holds. Then, the if-statement condition of

line 56

IJESPR www.ijesonline.com

A	Algorithm 6: Self-stabilizing replication for guards and leader, code for cloudlet p_i
112 I	Interfaces: <i>fetch()</i> next multicast message, <i>apply(state, msg)</i> applies the step <i>msg</i> to <i>state</i> (while producing side effects), <i>synchState(replica)</i> returns a replica consolidated state, <i>synchMsgs(replica)</i> returns a consolidated arrayof last delivered messages, <i>failureDetector()</i> returns a vector of processor ids, <i>cloudID</i> returns the address of the Cloud;
113	Variables: $rep[] = vie(v = ID, set, statu) \in \{Propose, Install, Multicast, \}(multicast round number) rnd, (replica) state, (last delivered messages) msg[n] (to the state machine), (last fetched) input (to the state machine), propV = \langle ID, set \rangle, (recently live and connected component) FD): an array of the state machine's replica, where rep[i] refers to the one that processor p_i maintains, and rep[j] refers to the last arriving message from p_j containing p_j's rep[j]. FD stores the failureDetector() output, i.e., the set of processors that the failure detector considers as active. myLeader stores the id of the local leader; if none. The view.ID (and propV.ID) is composed by the id and leader sequence installing the view, and counter cnt, in case the same leader installs a new view;$
114 S	Shared registers: info and data: as in Algorithm 3;
115] 116	Macros: $roundProceedReady() = \{ (\forall p_j \in view.set: rep[j].(view, status, rnd) = (view, status, rnd)) \lor ((status /= Multicast) \land [(\forall p_j \in propV.set : rep[j].(propV, status) = (propV, Propose)) \lor (\forall p_j \in propV.set : rep[j].(propV, status) = (propV, Install))] \}; coordinatePropose() = \{(state, msg, status) \leftarrow (synchState(rep), synchMsgs(rep), install);$
117 (118) 119 , 120 <i>f</i>	$coordinateInstall() = \{(view, status, rna) \leftarrow (prop \lor, Multicast, 0)\};$ $roundReadyToFollow() = \{rep[myLeader].rnd = 0 \lor rnd < rep[myLeader].rnd \lor rep[myLeader].(view /= prop \lor)\};$ $followPropose() = \{(status, prop \lor) \leftarrow rep[myLeader](status, prop \lor)\};$ $followInstall() = \{rep[i] \leftarrow rep[myLeader]\};$
121	followMcastRnd() {rep[i] \leftarrow rep[myLeader]; apply(state, rep[myLeader].msg); input \leftarrow fetch();}
122 L	procedure coordinateMcastRnd() do begin
123	$apply(state, msg); input \leftarrow fetch();$
124	for each $p_j \in C$ do if $p_j \in view.set$ then $msg[j] \leftarrow rep[j].input$ else $msg[j] \leftarrow \bot$;
125	write(data, (i, lLeader, rnd, rep[i].state));
126	_ rnd \leftarrow rnd + 1; if rnd = MAXINT then view.set $\leftarrow \perp /*$ Forces a view change in line 131/;
127	do forever /* use predefined periodicity */ begin
127	$FD \leftarrow failureDetector()$
129	let (IDevices, [Cloudlets, ILeader, IGuards) := read(info);
130	if <i>lleader.id = i</i> myLeader = i then (status, propV, myLeader) \leftarrow (Propose, { <i>lLeader.cnt = 0</i>), FD \cap
	$(Guards \cup \{i\})\rangle, i);$
131	if $lLeader.id = i$ myLeader = i ((status = Multicast view.set = S)) (k tatus \Rightarrow Multicast propV.set \approx S)))
132	then (status, propV, myLeader) \leftarrow (Propose, $\langle lLeader, cnt++, S \rangle$, $i \rangle$) where $S := FD$ (lGuards $\cup \{i\}$; if $k \neq i$ $\land i \in lGuards$ $\land k \in FD$, where $k = lLeader.id$ then (myLeader, status) \leftarrow (k, rep[k].status);
133	if <i>lLeader.id</i> = <i>i moundProceedReady()</i> then
134	if status = Multicast then coordinateMcastRnd();
135	else if status = Install then coordinateInstall();
136	_ else if <i>status</i> = <i>Propose</i> then <i>coordinatePropose</i> ();
137	else if $lLeader.id = i$ $i \land \in lGuards$ $kLeader.id \in FD \land roundReadyToFollow()$ then
138	if status = Multicast then followMcastRnd();
139	else if status = Install then followInstall();
140	_ else if <i>status</i> = <i>Propose</i> then <i>followPropose</i> ();
141	if <i>lLeader.id</i> $= i \land i \in lGuards \land lLeader.id \notin FD$ then
	$my_{Leader} \leftarrow \pm;$ while (aata, (i, i.e., aaer, ma, rep[1], state));
142	ense in myLeauer $l = \pm$ then send rep[1] to myLeauer; if $ _{aadar id} = i$ then $\forall k \in G_{uards} \cap FD$ send $ _{ran[i]}$ to p_i :
143	if $cret = MAYINT$ then $sond(cloudID / PESET)):$
144	I CHI = IMAAHVI then senu(CIOHUHD, (RESE1));
145	upon message <i>m</i> arrival from p_j up $rep(j) \leftarrow m$;
do Th	es not hold again during R , and, within O(1) cycles, the if-statement condition of line 48 holds. us, the system reaches R within O(1) cycles.

Evaluation 6

The previous section details the correctness proof of our self-stabilizing algorithmic process which, in contrast to the current state-of-the-art, shows that even in the presence of failures a fog ecosystem can always recoverin constant time and compute analytic insights from IoT data.



Figure 5: High-Level Overview of the Bus Network Topology

This section introduces a comprehensive evaluation of the effectiveness and runtime overhead of our framework. First, we measure *information delay* that is the time interval required for IoT data to be propagated in the network for analytics to be correctly derived in the presence of multiple and different failures (*e.g.*, cloudlet fail-stop, communication link drops). Second, we measure the additional *runtime footprint* that our framework incurs to exemplary state-of-the-art distributed control planes (*e.g.*, istio). This provides a detailed overview of what is the cost, in terms of network overhead, of maintaining *datafreshness* and analytic computation *correctness* in the presence of failures. *Results show that with our self-stabilizing framework, control planes are able to compute analytics correctly with the information delay*

maintained relatively stable despite of the presence of failures, while the network overhead scales linearlytowards the IoT load, as required by O1-O4 (Section 3).

For the experimentation, we introduce a real-world use-case of a smart city Bus Network Service (BNS) evaluated under various execution scenarios. We opt to focus on experiments that use a publicly available and real-world workload to truly reveal the strengths of our framework and its ability to deal with high workload. Specifically, the workload originates from the Dublin smart city Bus Network Service [2], comprised of 40GB of compressed data, tracking for 1 month the bus routes of 968 buses (Jan. 2013). Each bus is equipped witha GPS tracking device recording every 1s location coordinates and the current bus route delay. Figure 5 depicts a high-level overview of the BNS topology, where 16 cloudlets are deployed across Dublin's majorcity regions, denoted for clarity as A_x , to decentralize the BNS and increase the system responsiveness. We note that a bus route may span across different city regions and a bus can be connected to multiple cloudlets depending on the cloudlet coverage. Each cloudlet serves as an analytics engine that aggregates local bus updates and propagates an alert to traffic operators (central cloud service) when 10 or more buses in a city area are reporting, in a 5min sliding window, delays over one standard deviation from the previous weekly mean.

To experiment with large-scale deployments and ensure both result reproducability and algorithm adop- tion, we have designed a simulation testbed inspired by Kompics [6], an open-source distributed systems message-passing component model, and extended the entity behavior model to facilitate fault models over a

control plane in fog computing ecosystems². The testbed is run in an Openstack private cloud on a server configured with 16VCPU clocked at 2.66GHz, 16GB RAM and 260GB disk. The network configuration between testbed entities adopts a gaussian kernel with the following mean values: (i) Cloudlet-to-Cloud latency 100ms; (ii) IoT-to-Cloud latency 250ms; (iii) intra-region Cloudlet-to-Cloudlet latency 10ms; (iv) inter-region Cloudlet-to-Cloudlet latency 100ms; and (v) IoT-to-Cloudlet latency 20ms. We opt for these specific capabilities so that the testbed resembles an actual geo-distributed fog deployment over a city envi- ronment. All simulation scenarios are run 100 times with cloudlets and IoT devices starting at randomized time intervals. For the IoT device placement, we have implemented the registration interface of Algorithm 2so that when an IoT device (*e.g.*, a bus) requests to join the network, the central authority (*e.g.*, the cloud) responds with a list of valid cloudlets that are the "closest" to the device in the device's operating (city) region. The same strategy will hold for when the device has changed it's operating region (*e.g.*, bus moves from A_1 to A_2). Finally, the selection of the leader and the guards was done randomly, since our cloudletsare homogeneous.

For the widespread experimentation of different fault scenarios over the testbed, we adopt the NetflixChaos Monkey framework [1]. This enables the configuration and (random) selection of faults and entities to infest at given time intervals, or at random, depending on the evaluation scenario. Unless otherwise stated, the aforementioned topology and network configuration will be considered as the *baseline configuration*.

Information Delay

In this set of experiments, we show the effect of different failures to the timeliness of analytic computation. We consider four experiment runs with faults injected at random and examine how information delay is affected by:

- randomly failing a different number of regular cloudlets;
- failing the guards;
- failing the leader;
- randomly dropping the communication link between IoT devices and cloudlets.

Figure 6 depicts the information delay as the number of concurrently failing cloudlets increases. In this box-plot the median information delay is denoted by the line in the box, while the box length extends betweenthe first and third quantile with outliers depicted as independent points. With zero cloudlets we denote theinformation delay in normal operation (without failures). From Figure 6, we observe that information delayis not affected, despite slight deviations, while the number of failing cloudlets remains under 7. After this, randomly selecting concurrent cloudlets hinders the extreme case of wiping out all cloudlets of a city region. This results in added delay as IoT data for the specific region must be directly propagated to the cloud. For this experiment run, system recovery is only required when an IoT device is left with no cloudlet in itscoverage. In this extreme case, the IoT device must contact the cloud to validate the registration. However, the involvement of the cloud naturally hinders a communication overhead. Thus, despite information delaysfor extreme cases of concurrent cloudlet failures, analytic computation is correct at all times while the systemrecovers from faults in a bounded number of communication rounds, as required by O1 and O3 (Section 3). Figure 7 depicts how information delay is affected by the failure of the control plane guards and leader when the baseline deployment is configured with two guards. We observe that the timeliness of analytics computation is neither affected by the failure of the leader or the guards. This concurs with the correctnessproof that shows that, the self-stabilizing fog ecosystem can return back to a legal state within (1) *time*,

which is sufficient to propagate information without delay, as required by O4 (Section 3).

The next experiment run studies how information delay is affected by the temporary **(D**op of the networklink between IoT devices and cloudlets. To achieve this, we artificially block for a

predefined interval the

IJESPR www.ijesonline.com





link between affected IoTs and cloudlets in each region, thus maintaining only the link with the cloud. In Figure 8 we observe that the information delay increases as more devices experience a link drop. This occurs because the affected IoTs detect the link absence and, thus, must communicate with the cloud for updates which takes more time. Still, *analytics are computed without corrupted or missing IoT data*. This extreme case, of failing all the communication links among IoT and cloudlets, highlights the importance of having a sufficient amount of cloudlets in each region to cope with concurrent link failures.

Runtime Footprint

In this set of experiments, we provide an analysis depicting the network overhead of different components comprising our framework and the experiment testbed. Figure 9 depicts the network traffic over the data and control plane for a simulation run of the baseline configuration when random failures of the cloudlets' leader, guards, and cloudlets are introduced. The figure depicts the network overhead for 5min where the 30sbootstrap period is omitted. First, we observe four distinct segments (separated by vertical lines). During each segment our framework maintains a stable message exchange rate for both planes, with the data plane traffic approximately x3.5 higher than the control plane traffic. In the first segment (30s to 75s) the system exhibits no faults. At the 75th second, the leader fails and we observe a slight drop in both the control planetraffic (from 950KB/s to 850KB/s) and the data plane (from 3300KB/s to 3100KB/s). When the cloud discovers the leader failure, it elects a new leader at the 88th second and the system recovers back to a legal state, with a slight increase of the control plane traffic (900KB/s). Next, at the 150th second the two guards fail and the control plane traffic falls to 700KB/s while the data plane traffic falls to 2600KB/s. As before, the cloud elects two new guards and the control plane traffic stabilizes at 750KB/s. Finally, at the 225th second (4th segment) three cloudlets fail and both control and data traffic drop to 550KB/s and 2100KB/s, respectively. These results show that a constant number of messages is exchanged, validating the objectivesO1 and O3 (Section 3).

Next, we show that the control and data plane network traffic scales linearly towards the number of different system entities, as required by O2. Table 1 shows the results of different configurations in percentage increments from the baseline.



Figure 7: Information delay vs number of concurrent guard and leader failures

Guards. We observe that the overhead of adding guards increases linearly. Specifically, each additional guard adds an overhead in the range of 4.75 5.68% for the control plane traffic and 4.82 5.02% for the data plane traffic. It is worth pointing out that the previous experiment in Figure 7 showed that even with all the guards failing concurrently, the information latency remains stable, and therefore, for the studied baseline configuration, having two guards balances well the trade-off between overhead and information delay.

Cloudlets. The overhead of adding extra cloudlets, for redundancy purposes, scales linearly while the IoT load remains stable. Specifically, each additional cloudlet adds an overhead in the range of $7.26 \quad 9.03\%$ for the control plane, and for the data plane the increment is approximately 6.2%. Obviously, the trade-off is straightforward. Increasing the cloudlets, decreases the probability of delaying information propagation for a city region, *e.g.*, as in the case of Figure 6 after 7 cloudlets, at the cost of higher network traffic.

IoTs. By increasing the workload (IoT devices), again, the network overhead is linearly increased. Each additional IoT device adds a 0.094% overhead on the control plane traffic, while for the data plane the incre-ment ranges between 0.085 0.087%. This increase is attributed to the fact that each cloudlet communicates with more IoT devices.

7 Conclusions

In this paper we introduced a fault-tolerant framework for distributed control planes that enables fog services to cope with a very broad fault model. To this end, we presented self-stabilizing algorithms that guarantee automatic recovery within a constant number of communication rounds without the need for external (hu- man) intervention. Using real-world data and actual queries of interest from an intelligent transportation service, we demonstrate the performance gains of our framework, and thus the promise of self-stabilization in fog computing. Our results show that despite information delays for extreme cases of concurrent cloudlet fail- ures, analytic computation is correct, while the network overhead is proportional to the number of cloudlets, guards, and devices. We believe that our self-stabilizing framework is applicable to a wide range of fog services requiring strong fault-tolerance guarantees.



Figure 8: Information delay vs concurrent fail-stop IoT-Cloudlet network links

References

- [1] Chaos monkey tool.
- [2] Dublin smart city bus network data.
- [3] Istio.
- [4] Linkerd.
- [5] Abdulaziz Alarifi, Fathi Abdelsamie, and Mohammed Amoon. A fault-tolerant aware scheduling method for fog-cloud environments. *PLOS ONE*, 14(10):1–24, 10 2019.
- [6] Cosmin Arad, Jim Dowling, and Seif Haridi. Message-passing concurrency for scalable, stateful, recon- figurable middleware. In *Middleware 2012*, pages 208–228. Springer Berlin Heidelberg, 2012.
- [7] Alexander Binun, Mark Bloch, Shlomi Dolev, Ramzi Martin Kahil, Boaz Menuhin, Reuven Yagel, Thierry Coupaye, Marc Lacoste, and Aurélien Wailly. Self-stabilizing virtual machine hypervisor archi- tecture for resilient cloud. In 2014 IEEE World Congress on Services, pages 200–207, 2014.
- [8] Luiz Bittencourt, Roger Immich, Rizos Sakellariou, Nelson Fonseca, Edmundo Madeira, Marilia Curado, Leandro Villas, Luiz DaSilva, Craig Lee, and Omer Rana. The internet of things, fog and cloud continuum: Integration and challenges. *Internet of Things*, 3-4:134 155, 2018.

- [9] Peva Blanchard, Shlomi Dolev, Joffroy Beauquier, and Sylvie Delaët. Practically self-stabilizing paxos replicated state-machine. In *Networked Systems Second International Conference, NETYS 2014, Marrakech, Morocco, May 15-17, 2014. Revised Selected Papers*, pages 99–121, 2014.
- [10] Marco Canini, Iosif Salem, Liron Schiff, Elad Michael Schiller, and Stefan Schmid. Renaissance: A selfstabilizing distributed SDN control plane. In *ICDCS 2018, Vienna, Austria, July 2-6, 2018*, pages233–243, 2018.
- [11] S. Chattopadhyay, S. Chatterjee, S. Nandi, and S. Chakraborty. Aloe: An elastic auto-scaled and selfstabilized orchestration framework for iot applications. In *IEEE INFOCOM*, pages 802–810, 2019.
- [12] Edsger W. Dijkstra. Self-stabilizing systems in spite of distributed control. Commun. ACM, 17(11):643-644, 1974.
- [13] Shlomi Dolev. Self-Stabilization. MIT Press, 2000.
- [14] Shlomi Dolev, Chryssis Georgiou, Ioannis Marcoullis, and Elad Michael Schiller. Practically-selfstabilizing virtual synchrony. J. Comput. Syst. Sci., 96:50–73, 2018.
- [15] Michael J. Fischer and Michael Merritt. Appraising two decades of distributed computing theory research. *Distributed Computing*, 16(2-3):239–247, 2003.
- [16] Yotam Harchol, Aisha Mushtaq, James McCauley, Aurojit Panda, and Scott Shenker. CESSNA: Resilient edge-computing. In Proceedings of the 2018 Workshop on Mobile Edge Communications, pages 1–6. ACM, 2018.
- [17] T. He, H. Khamfroush, S. Wang, T. La Porta, and S. Stein. It's hard to share: Joint service placement and request scheduling in edge clouds with sharable and non-sharable resources. In *ICDCS 2018, Vienna, Austria, July 2-6, 2018*, pages 365–375, 2018.
- [18] Cheol-Ho Hong and Blesson Varghese. Resource management in fog/edge computing: A survey on architectures, infrastructure, and algorithms. *ACM Comput. Surv.*, 52(5), September 2019.
- [19] Radia Perlman. Interconnections (2Nd Ed.): Bridges, Routers, Switches, and Internetworking Protocols. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2000.
- [20] Justine Sherry, Peter Xiang Gao, Soumya Basu, Aurojit Panda, Arvind Krishnamurthy, Christian Maciocco, Maziar Manesh, João Martins, Sylvia Ratnasamy, Luigi Rizzo, and Scott Shenker. Rollback- recovery for middleboxes. SIGCOMM Comput. Commun. Rev., 45(4):227–240, August 2015.
- [21] W. Shi and S. Dustdar. The promise of edge computing. Computer, 49(5):78–81, May 2016.
- [22] Gerry Siegemund and Volker Turau. A self-stabilizing publish/subscribe middleware for iot applications. *ACM Trans. Cyber-Phys. Syst.*, 2(2):12:1–12:26, June 2018.
- [23] Olena Skarlat, Matteo Nardelli, Stefan Schulte, Michael Borkowski, and Philipp Leitner. Optimized iot service placement in the fog. *Service Oriented Computing and Applications*, 11(4):427–443, 2017.
- [24] D. Trihinas, L. Chiroque, G. Pallis, A. Fernandez, and M.D. Dikaiakos. ATMON: Adapting the "Temporality" in Large-Scale Dynamic Networks. In 38th IEEE International Conference on Distributed Computing Systems (ICDCS 2018), 2018.
- [25] D. Trihinas, G. Pallis, and M. Dikaiakos. Low-cost adaptive monitoring techniques for the internet of things. *IEEE Transactions on Services Computing*, pages 1–1, 2018.
- [26] Demetris Trihinas, George Pallis, and Marios Dikaiakos. ADMin: adaptive monitoring

IJESPR www.ijesonline.com

dissemination for the internet of things. In IEEE INFOCOM 2017, Atlanta, USA, May 2017.

- [27] Chao Wang, Christopher Gill, and Chenyang Lu. Frame: Fault tolerant and real-time messaging for edge computing. In *ICDCS*, pages 976–986, 2019.
- [28] K. Wang, Y. Shao, L. Xie, J. Wu, and S. Guo. Adaptive and fault-tolerant data processing in healthcare iot based on fog computing. *IEEE Transactions on Network Science and Engineering*, pages 1–1, 2019.