

Towards Self-Adaptive Peer-to-Peer Monitoring for Fog Environments

Vera Colombo
v.colombo51@campus.unimib.it
University of Milano-Bicocca
Milan, Italy, IT

Michele Ciavotta
michele.ciavotta@unimib.it
University of Milano-Bicocca
Milan, Italy, IT

Alessandro Tundo
alessandro.tundo@unimib.it
University of Milano-Bicocca
Milan, Italy, IT

Leonardo Mariani
leonardo.mariani@unimib.it
University of Milano-Bicocca
Milan, Italy, IT

ABSTRACT

Monitoring is a critical component in fog environments: it promptly provides insights about the behavior of systems, reveals Service Level Agreements (SLAs) violations, enables the autonomous orchestration of services and platforms, calls for the intervention of operators, and triggers self-healing actions.

In such environments, monitoring solutions have to cope with the heterogeneity of the devices and platforms present in the Fog, the limited resources available at the edge of the network, and the high dynamism of the whole Cloud-to-Thing continuum.

This paper addresses the challenge of accurately and efficiently monitoring the Fog with a self-adaptive peer-to-peer (P2P) monitoring solution that can opportunistically adjust its behavior according to the collected data exploiting a lightweight rule-based expert system.

Empirical results show that adaptation can improve monitoring accuracy, while reducing network and power consumption at the cost of higher memory consumption.

CCS CONCEPTS

• **Computer systems organization** → **Self-organizing autonomous computing**; *Peer-to-peer architectures*; • **Information systems** → *Expert systems*.

KEYWORDS

self-adaptive, peer-to-peer, monitoring, fog computing

1 INTRODUCTION

Fog computing has been proposed to fill the gap between the Cloud and the Internet-of-Things (IoT), providing a continuum of computing, storage, and networking facilities designed to address emerging scenarios related to several domains, such as smart cities, transportation, industry 4.0, and e-health [7, 16, 41]. Its infrastructure is made of a large number of widely distributed heterogeneous nodes, which can include network devices (e.g., routers, gateways, access points, and base stations), micro-clouds and computing servers [16, 41].

Monitoring is an intrinsically pivotal component of the Fog as it is the entity responsible for providing reliable information about the behavior of infrastructure and services [1, 2, 32]. Monitoring data are used in several tasks, including the autonomous orchestration

of services and platforms [32], the detection of Service Level Agreements (SLAs) violations [4], the activation of operator-mediated procedures [6], and self-healing actions [41].

Although a large number of monitoring solutions are already available for the Cloud [5, 9, 10, 23, 25, 29, 35, 37, 38], they are well-known to badly adapt to heterogeneous and massively distributed infrastructures characterized by frequent changes to the topology of the nodes, such as the Fog [41]. For instance, Abderrahim et al. [1] discuss how a fog monitoring system, unlike cloud-specific solutions, must make particularly good use of the available resources, must be resilient to changes in the topology (e.g., nodes joining and leaving the network) and network conditions (e.g., communication links may not always be fully operational).

Peer-to-Peer (P2P) architectures have been recently investigated as viable approaches to effectively address monitoring in the Fog [1, 13, 14, 41]. Indeed, as they are “self-organizing systems of equal, autonomous entities (peers) which aim for the shared usage of distributed resources in a networked environment avoiding central services” [24], they represent a legitimate option to address the dynamism of the Fog without imposing strong constraints on the stability of the target environment. For instance, FogMon [13] is a P2P monitoring solution targeting fog infrastructures that can monitor hardware resources, end-to-end network QoS, and connected IoT devices. SkyEye [14] is a P2P tree-based monitoring approach that provides lightweight continuous monitoring of all peers in its network.

Unfortunately, although these monitoring systems show some degree of adaptivity thanks to the features provided by P2P architectures (e.g., they can tolerate node disconnections and broken communication links), they lack adaptation mechanisms that take into account the monitored indicators [1, 13, 14]. In fact, the collected indicators reveal important information about the monitored system and its environment, and can be exploited to increase the awareness and adaptability of the monitoring system itself. For example, a monitoring node running in a device exhausting its battery may stop monitoring the non-essential indicators. Similarly, the trend of a monitored indicator can be used to optimize the sampling rate to avoid wasting resources (e.g., increasing/decreasing the sampling rate based on the degree of stability of the indicator).

In this paper, we address this limitation by proposing a *self-adaptive P2P monitoring* solution that exploits a hierarchical P2P architecture and embeds adaptive behaviors defined according to

the MAPE-K feedback loop [19]. The proposed monitoring system can abstract the monitored indicators and activate countermeasures based on the status of these indicators. Countermeasures are defined by exploiting a lightweight rule-based system embedded in the peers. In the paper, we show how this framework can be used to activate two example countermeasures that can dynamically change the set of monitored indicators and their sampling rate.

We implemented our solution as an extension of FogMon [13], a state-of-art P2P monitoring solution for the Fog, and released it publicly¹. The empirical evaluation of the accuracy and effectiveness of the adaptive version of the monitoring solution in contrast with the non-adaptive one shows that adaptive behaviors can be used to increase the accuracy of the collected data and save network and power consumption at the cost of higher memory consumption. This is an interesting trade-off since devices are often reasonably equipped in terms of memory, while network bandwidth and power are critical resources at the edge [1].

In a nutshell, the main contributions of this paper are:

- A self-adaptive monitoring solution that modifies its behavior by running countermeasures defined over an abstraction of the monitored indicators.
- The instantiation of the MAPE-K loop on the case of P2P fog monitoring.
- The definition of two countermeasures to self-adapt the monitoring system: (i) *Change Rate*, which can be used to dynamically alter the data collection rate; and (ii) *Select Indicators*, which can be used to dynamically control the activation and deactivation of the probes.
- A publicly available prototype obtained by extending the FogMon P2P monitoring system.
- An empirical evaluation of the monitoring accuracy and efficiency of the adaptive solution in comparison to the non-adaptive counterpart.

The rest of this paper is organized as follows. Section 2 discusses the use of P2P architectures for monitoring in the Fog. Section 3 provides a rigorous definition of the self-adaptive P2P monitoring solution proposed in this paper. Section 4 briefly describes our prototype implementation. Section 5 presents the empirical results about monitoring accuracy and efficiency. Section 6 discusses related work. Finally, Section 7 provides final remarks.

2 P2P MONITORING

A P2P architecture consists of a network of autonomous self-organizing entities (i.e., peers) that employ distributed resources to accomplish a common task in a decentralized fashion, thus, without relying on central services [22, 24, 31].

The P2P architecture provides the applications with the capability to deal with some of the highly dynamic traits of fog computing, increasing the tolerance to both network failures and nodes joining and leaving the system [1, 11, 13]. Furthermore, it provides autonomy, scale, and robustness, which are critical capabilities to operate in such an environment [34]. Finally, P2P architectures make monitoring data available across the network without relying on a single centralized component, but rather on a set of peers constituting a self-organized overlay network. This is particularly

¹<https://github.com/veracoo/FogMon/tree/adaptive-fogmon>




Figure 1: Hierarchical P2P monitoring architecture proposed by Forti et al. [13].

beneficial when the connectivity to the Cloud is limited, such as during disasters or severe network outages [41].

In this work, we use the two-tier hierarchical P2P monitoring architecture [13, 20, 40] proposed by Forti et al. [13] that is shown in Figure 1, with *Followers* at the lower tier and *Leaders* at the higher tier.

The benefit of employing such an architecture in the Fog is twofold. First, it implies different roles for peers running in different tiers, depending on the available resources. Followers are designed to collect data by running on the edge, within nodes and devices with limited resources. Leaders are designed to consume more resources to store the data received from the Followers while creating and operating the P2P network. Followers are connected to a single Leader and work in a classic client-server fashion [40]. These distinct roles can be used to opportunistically exploit the available resources, including the possibility to adapt to changing conditions (e.g., bandwidth or resource degradation) through dynamic peer promotion/demotion.

Second, it helps reducing the network overhead by limiting the amount of data transferred between the peers. Actually, Followers can forward data to their Leader only, leaving the thinner upper-tier with the responsibility of building a global state of the monitored resources by exchanging monitoring data among Leaders.

We refer to FogMon as reference implementation for this architecture [13]. In FogMon, the Followers monitor their own deployment node by probing hardware resources (i.e., CPU, memory, and hard disk), collecting end-to-end network QoS data (i.e., latency and bandwidth), and detecting IoT devices. Data is collected and sent to Leader nodes at a *fixed rate*. To limit network overhead, Followers send differential updates, that is, they only send data whose average or variance differ more than a *sensitivity* threshold (i.e., 10% by default) from the last value sent [13]. Leaders periodically aggregate monitoring data received from Followers, and share the aggregated data with the other Leaders through a gossip protocol [17].

3 SELF-ADAPTIVE P2P MONITORING

This section describes how P2P monitoring can be enhanced with self-adaptive capabilities to both make a better use of the available resources and enable the capability to promptly react to run-time




Figure 2: Monitor, Analyze, Plan, Execute, and Knowledge (MAPE-K) loop as proposed by Kephart and Chess [19].

events. We refer to the self-adaptive version of the P2P monitoring solution presented in this paper as *ADAPTIVEMON*, in contrast with the non-adaptive version that we refer to as *STATICMON*.

ADAPTIVEMON enriches the capabilities of the monitoring system by embedding the Monitor, Analyze, Plan, Execute, and Knowledge (MAPE-K) control framework [19], shown in Figure 2, within each peer. The *monitor* component of the MAPE-K loop collects data about a managed resource through *sensors*. In our case, this is achieved by the monitoring probes running within the peers. The *analyze* and *plan* steps analyze the collected data and plan for the countermeasures to be activated. Finally, the *execute* step exploits *effectors* to run the selected countermeasures. In our case, the countermeasures reconfigure the monitoring systems according to the observations collected from the managed resource. The *knowledge* about the managed resource is shared among all the components. Note that the MAPE-K loop runs within each peer, independently of the overall architecture, which gives peers the capability to run self-adaptive behaviors regardless of their role within the architecture.

In the following, we describe how the components of the MAPE-K loop embedded in the peers are defined, and present two countermeasures that we experienced in our prototype implementation, namely (i) *Change Rate*, which adjusts the rate Followers sample and forward data to their Leader, and (ii) *Select Indicators*, which dynamically activates and deactivates the set of monitored indicators.

3.1 Knowledge

The knowledge exploited in *ADAPTIVEMON* consists of the monitored *indicators*, which represent the raw knowledge about the monitored resource, and the associated *logical states*, which capture the semantics of the values of an indicator.

Definition 3.1 (Indicator Values and time series). Given a monitored indicator I and a domain D of values for I , $v_t^I \in D$ denotes the value of the indicator I at time t . A sequence of values for a same indicator generates a time series, that is, $v_i^I, v_{i+1}^I, \dots, v_k^I$ is a time series for indicator I .

Definition 3.2 (Logical States). Given a monitored indicator I and a finite set of abstract states S , $S_t^I \subseteq S$ represents a potentially empty

set of logical states associated with the indicator I at time t . The sequence of states sets $S_t^I, S_{t+1}^I, \dots, S_k^I$ associated with an indicator also forms a time series.

For sake of notation, we omit I when the indicator is obvious from the context.

While time series of values simply reflect the sequence of probed samples, the corresponding time series of logical states captures the state of an indicator at a specific time, revealing information about the monitored resource. For instance, an indicator might be unstable, too high, or within a normal range of values. These states can be derived from the time series of raw values and used to fire countermeasures, as explained below.

We defined a set of logical states useful for our countermeasures, but this set can be clearly extended depending on the countermeasures to be implemented. We summarize the rigorous definitions of the logical states in Table 1, while we discuss them informally below.

In case the domain of a metric is *categorical*, we defined two logical states: *stable* and *unstable*. A categorical indicator is stable at a time t if its value has been constant in the recent history of the execution, otherwise it is unstable.

In case the domain of an indicator is *numerical*, we defined seven states: *stable*, *unstable*, *normal*, *high*, *low*, *too high*, *too low*. A numerical indicator is stable at a time t if its value is close to its value at time $t - 1$, and the indicator had small variations in the recent history of the execution. Similarly, a numerical indicator is unstable at a time t if its value differs significantly from the value at time $t - 1$, and the indicator had significant variations in the recent history of the execution.

The remaining five states detect values that steadily stay in a given region of the domain. In particular, an indicator is too high or high if its value is above given thresholds and the indicator has been mostly above those thresholds in its recent execution history. Similarly, an indicator is too low or low if its value is below given thresholds and the indicator has been mostly below those thresholds in its recent execution history. Finally, an indicator is normal if its value is in the normal range, and the indicator has been mostly normal in its recent history of the execution. The thresholds for the various levels are defined on a per-indicator basis since they depend on both the indicator and the application domain. For instance, threshold values for memory and CPU consumption are clearly different, and threshold values of memory consumption for a memory-intensive application and a lightweight application are also different.

The set of the collected indicators, along with their raw values and state values represent the knowledge available to *ADAPTIVEMON*.

3.2 Monitor

Monitoring is rather natural and inexpensive in *ADAPTIVEMON* since Followers collect data from a monitored resource by construction, and thus the same data sent to Leaders is also available to the MAPE-K loop. If needed, extra indicators can be collected for the only purpose of controlling the adaptive behavior of the peers, even if not needed by the applications accessing the data produced by the monitoring system.

Indicator Type	State	Definition
Categorical	stable	$s_t = \text{stable} \iff v_t = v_{t-x} \quad \forall x \in [1, k]$
	unstable	$s_t = \text{unstable} \iff \exists x \in [1, k] \mid v_t \neq v_{t-x}$
Numerical	stable	$s_t = \text{stable} \iff Stab_t \geq p \cdot k \wedge v_t - v_{t-1} \leq \Delta_{max},$ $Stab_t = \{v_x - v_{x-1} \mid \leq \Delta_{max}\}_{x \in [1, k]}$
	unstable	$s_t = \text{unstable} \iff Stab_t < p \cdot k \vee v_t - v_{t-1} > \Delta_{max},$ $Stab_t = \{v_x - v_{x-1} \mid < \Delta_{max}\}_{x \in [1, k]}$
	too high	$s_t = \text{too high} \iff Too_High_t \geq p \cdot k \wedge v_t \in I_{too_high},$ $Too_High_t = \{v_x \in I_{too_high}\}_{x \in [0, k]}$
	high	$s_t = \text{high} \iff High_t \geq p \cdot k \wedge v_t \in I_{high},$ $High_t = \{v_x \in I_{high}\}_{x \in [0, k]}$
	normal	$s_t = \text{normal} \iff Normal_t \geq p \cdot k \wedge v_t \in I_{normal},$ $Normal_t = \{v_x \in I_{normal}\}_{x \in [0, k]}$
	low	$s_t = \text{low} \iff Low_t \geq p \cdot k \wedge v_t \in I_{low},$ $Low_t = \{v_x \in I_{low}\}_{x \in [0, k]}$
	too low	$s_t = \text{too low} \iff Too_Low_t \geq p \cdot k \wedge v_t \in I_{too_low},$ $Too_Low_t = \{v_x \in I_{too_low}\}_{x \in [0, k]}$

Symbols	Definition
$ $	Cardinality of a set.
k	Number of samples considered in the recent history of an indicator.
$p \in [0, 1]$	Tolerance parameter that indicates the percentage of k samples. It must satisfy the constraint that characterizes the state definition.
Δ_{max}	Maximum delta allowed to consider an indicator as stable.
$I_{too_high} = [too_high, +\infty)$	Interval of indicator values considered <i>too high</i> .
$I_{high} = [high, too_high)$	Interval of indicator values considered <i>high</i> .
$I_{normal} = (low, high)$	Interval of indicator values considered <i>normal</i> .
$I_{low} = (too_low, low]$	Interval of indicator values considered <i>low</i> .
$I_{too_low} = (-\infty, too_low]$	Interval of indicator values considered <i>too low</i> .

Table 1: States definitions for categorical and numerical indicators.

The monitoring behavior is controlled by a *sampling rate* parameter that determines how frequently values v_t^I are collected and forwarded to Leader peers.

3.3 Analyze

The analysis mainly consists of a data processing routine that converts the raw values collected for every indicator into its logical state representation. In particular, the analysis process accesses the collected values and applies the definitions reported in Table 1 to generate a time series of logical states for every monitored indicator.

Figure 3 visually exemplifies the logical states that can be associated with a time series, according to the definitions in Table 1. The logical states are represented as annotations on the X-axis. Note that in the example, depending on the shape of the curve, a same point may have 0, 1 or up to 2 logical states associated.




Figure 3: An example of the computed states with respect to the time series values at different time instants.

3.4 Plan

ADAPTIVEMON embeds the lightweight CLIPS expert system [26], which is responsible for determining the countermeasures that have to activate according to the accumulated knowledge based on a set of adaptation rules.

An adaptation rule consists of two parts: an *antecedent*, that is a set of conditions on the logical states of the indicators that must be satisfied to fire the rule, and a *consequent*, that is a countermeasure to be executed.

We extended CLIPS by implementing functions that can be used to evaluate the logical states of the indicators, such as functions that can check specific conditions on the last few samples of an indicator. These functions can be used as part of the adaptation rules specified using the CLIPS Domain-Specific Language (DSL). For instance, Listing 1 shows an adaptation rule defined to fire the *Change Rate* countermeasure when the CPU consumption has a stable trend. The symbol => separates the antecedent and the consequent of the rule.

It is worth noting that the example adaptation rule uses some of the functions we defined in ADAPTIVEMON. It checks if the CPU consumption is in a stable state with the `is_indicator_in_state` function and it computes the new rate for such indicator by executing the `compute_indicator_rate` function. The new rate value is in turn used, along with other variables retrieved from the knowledge base, by the `change_rate` function that implements the *Change Rate* countermeasure.

During the plan phase, ADAPTIVEMON uses the CLIPS expert system to take adaptation decisions, that is, CLIPS evaluates the antecedents of every rule and inserts the countermeasures activated by the consequent of the fired rules in a priority-based queue. Also, CLIPS handles the activation of the rules by preventing their simultaneous execution.

To demonstrate the self-adaptive capabilities of ADAPTIVEMON we defined two specific countermeasures: *Select Indicators* and *Change Rate*.

The *Select Indicators* countermeasure can change the set of monitored indicators, either activating or deactivating some of them. The *Change Rate* countermeasure changes the rate used to sample and send data to Leaders based on the current trends of the indicators. The goal of the countermeasure is to gradually increase (decrease) the rate while the monitored indicator is less (more) stable. In particular, the countermeasure updates the sampling rate of the indicator I within predetermined boundaries proportionally to the number of consecutive samples with the same logical state out of the last k samples collected.

We exploited these countermeasures in the context of several adaptation rules. For instance, we defined two rules that can enable/disable monitoring for every indicator different from power consumption if the power is above/below a given threshold, to limit the chance a battery-powered device is abruptly shut down. We also defined two of rules to adapt the sampling and forwarding rate of CPU indicator to its trend.

Listing 1: An example rule that uses the *Change Rate* countermeasure written with the CLIPS DSL. The symbol => separates the antecedent and the consequent of the rule. The salience value represents the rule priority. The bind operator assigns the result of a function call to a variable.

```
(defrule adapt_cpu_rate_when_stable (declare (salience
  10))
  (is_indicator_in_state (indicator cpu) (state stable))
  (has_parameter (rate ?current_rate))
  =>
  (bind ?num_of_states (count_indicator_states_in "cpu" "
    stable"))
  (change_rate "cpu" (compute_indicator_rate "stable" ?
    num_of_states ?current_rate))
  )
```

3.5 Execute

In this phase, ADAPTIVEMON merely executes countermeasures by running their implementation according to their priority of activation. The actual countermeasures we defined act on the configuration of the peers adapting their behavior to the evolution of the indicators. The actuation interface is straightforward since a peer can directly access the internal variables that govern its behavior.

4 PROTOTYPE

We implemented ADAPTIVEMON by extending the open-source C++ FogMon P2P monitoring tool [13] along multiple dimensions.

In particular, (i) we extended the structure of the peer's local storage (based on the SQLite² database) to store the logical states used to classify of the monitored indicators; (ii) we embedded the CLIPS rule-based expert system [26] to support the implementation of adaptation rules; (iii) we extended the peers to incorporate adaptive behaviors; (iv) we added helper functions that can be used as part of the adaptation rules to interact with the knowledge; and (v) we implemented the *Select Indicators* and *Change Rate* countermeasures to dynamically change the set of the collected indicators and the sampling rate parameters. We have not extended the set of monitored indicators, since the indicators already collected by FogMon to monitor the environment were already sufficient to control the activation of our countermeasures.

The resulting prototype is publicly available with an open-source licence at <https://github.com/veracoo/FogMon/tree/adaptive-fogmon>.

5 EVALUATION

To assess ADAPTIVEMON, we investigate the following two research questions about its effectiveness (monitoring accuracy) and efficiency (resource consumption).

5.1 Research Questions

RQ1 (Monitoring Accuracy) - Can ADAPTIVEMON improve the monitoring accuracy of STATICMON? This research question investigates whether the *Change Rate* policy of ADAPTIVEMON can provide a better monitoring accuracy than STATICMON, considering multiple representative trends of the monitored indicators.

²<https://www.sqlite.org>

RQ2 (Resource Consumption) - Can ADAPTIVEMON save node resources compared to STATICMON? This research question studies whether the adaptive behavior of ADAPTIVEMON reduces resource utilization compared to STATICMON. We assess the impact of the *Change Rate* and *Select Indicators* countermeasures on resource consumption, both individually and jointly.

We run all the experiments on a Linux virtual machine (Intel i7-9700 CPU @ 3.00GHz x 4, 4 GB RAM, 13 GB SSD, Ubuntu 20.04 LTS 64-bit, Docker v20.10.0). All the peers run inside dedicated Docker containers, deployed on the same host, and communicate over a bridged network. We limit the computational and network resources of the container executing the Follower agent to reproduce a scenario involving resource-constrained and battery-powered devices. Reference devices are single-board computers (SBC) and micro-controller units (MCU) [12, 18, 27]; thus, container resources are upper-bounded at 5% of one CPU core, 20 MB of RAM, and 1 Mbps of bandwidth.

We measure accuracy and resource consumption at the level of individual peers to obtain results that do not depend on the number of co-deployed peers. Thus, each experiment involves one Leader and one Follower (of ADAPTIVEMON or STATICMON, respectively). Cumulative resource consumption for multiple nodes can be derived by scaling the results proportionally.

5.2 RQ1 - Monitoring Accuracy

We investigate the accuracy of the collected data for both ADAPTIVEMON and STATICMON considering synthetic indicators following different representative trends. More in detail, we created a probe reporting readings from such indicators. This allowed us to test the correctness of ADAPTIVEMON’s adaptive behavior and verify its effectiveness. We defined 5 scenarios (also referred to as time series in the following) mimicking different key cases for an indicator conventionally ranging between 0 and 1:

- (1) *stable* is a time series representing a regular and stable, almost constant, trend. It is generated by alternating two close values (0.8 and 0.83), each of which remains stable for 14 seconds;
- (2) *unstable* is a time series that represents an irregular and erratic indicator with fluctuating values in the range [0.5, 0.85]; a real-life trace on CPU utilization was used as a base.
- (3) *stable-unstable* is a time series that alternates phases of stability with phases of instability, with each phase lasting for about 150 seconds;
- (4) *random* is a time series with chaotic and totally unpredictable values; it is generated by a sequence of random values uniformly distributed in the range [0, 1];
- (5) *spiky* is a time series with mostly regular values interleaved with rare spikes; it is generated by alternating stable values for 28 seconds, unstable values for 12 seconds, and then a spike value for 4 seconds.

Every time series has a duration of 10 minutes, except for the stable-unstable time series that lasts 20 minutes since it is a combination of the stable and unstable time series.

STATICMON and ADAPTIVEMON’s Follower peers are configured to forward the average of the last 20 measurements to their respective Leaders at each probing point. The STATICMON Follower probes a new value from the monitored metric at a fixed interval:

Scenario	Quality Metric	ADAPTIVEMON	STATICMON	Abs	Rel
<i>Stable</i>	RMSE (Follower)	0.019	0.020	- 0.181	- 5.0 %
	RMSE (Leader)	6.696	8.200	- 1.504	- 18.3 %
	Messages/second	0.027 m/s	0.040 m/s	- 0.013	- 32.5 %
<i>Unstable</i>	RMSE (Follower)	0.087	0.131	- 0.044	- 33.6 %
	RMSE (Leader)	2.428	5.033	- 2.605	- 51.7 %
	Messages/second	0.217 m/s	0.037 m/s	+ 0.180	+ 486.5 %
<i>Stable-unstable</i>	RMSE (Follower)	0.108	0.122	- 0.014	- 11.5 %
	RMSE (Leader)	5.269	6.546	- 1.277	- 19.5 %
	Messages/second	0.103 m/s	0.035 m/s	+ 0.068	+ 194.3 %
<i>Random</i>	RMSE (Follower)	0.235	0.321	- 0.086	- 26.8 %
	RMSE (Leader)	1.899	10.683	- 8.784	- 82.7 %
	Messages/second	0.217 m/s	0.037 m/s	+ 0.180	+ 486.5 %
<i>Spiky</i>	RMSE (Follower)	0.092	0.087	+ 0.005	+ 5.8 %
	RMSE (Leader)	5.713	6.251	- 0.538	- 8.6 %
	Messages/second	0.062 m/s	0.037 m/s	+ 0.025	+ 67.6 %
	Detected spikes	30 %	0	+30 %	-

Table 2: Accuracy of ADAPTIVEMON and STATICMON for the 5 scenarios. Green (Red) cells indicate a better (worse) result obtained by ADAPTIVEMON compared to the STATICMON.

every 30 seconds. ADAPTIVEMON exploits the *Change Rate* countermeasure to adjust the sampling rate to handle the variability of the monitored indicator. The hypothesis to test is that this can lead to improved monitoring accuracy because the Leader should have access to a higher number of samples when the monitored indicator is highly dynamic and fewer samples in the presence of more stable indicators.

We investigate the capability of the monitoring system to reconstruct the shape of the monitored indicators at the level of both the Follower, which directly samples the indicator, and the Leader, which collects a sequence of average values. We use the *Root Mean Square Error* (RMSE), which measures the differences between the original and the reconstructed indicator, as the primary quality metric. A smarter sampling strategy should achieve a lesser error. To appreciate the activity of the peers in relation to the monitored indicator, we also gauge the *messages/second* ratio, that is, the ratio of the messages sent by the Follower to the Leader. Finally, for the spiky time series, we also computed the *percentage of detected spikes*, which measures the capability of a monitoring technique to spot rare but significant events.

Results. Table 2 summarizes the results obtained by both ADAPTIVEMON and STATICMON for the considered 5 scenarios. The last two columns show the absolute (Abs) and relative (Rel) deviations between the ADAPTIVEMON and STATICMON results for any of the presented quality indicators. Green (Red, respectively) cells indicate a better (worst) result obtained by ADAPTIVEMON compared to the STATICMON baseline. We can observe that ADAPTIVEMON estimates the observed indicator more accurately than STATICMON at both levels of the Leader-Followers hierarchy in 4 out of 5 scenarios (*viz. stable, unstable, stable-unstable, random*). The reduction in the RMSE reached 33.6% at Follower level (*unstable* scenario) and 82.7% at Leader level (*random* scenario). A higher accuracy, however, comes at the cost of a higher number of messages exchanged

in the 4 scenarios where the monitored indicator is more erratic (nearly 5 times more than `STATICMON` in the worst case), and fewer messages produced when the indicator is stable (saving nearly one third of the messages). These results demonstrate the capability of `ADAPTIVEMON` to modulate resource consumption as needed, exchanging more messages only if the monitored indicator requires finer sampling, saving bandwidth otherwise.

Figures 4 and 5 exemplify the results of the comparison between `ADAPTIVEMON` and `STATICMON` for the *stable-unstable* scenario. We can observe that the time series reconstructed by `ADAPTIVEMON` (solid orange line) is closer to the reference indicator (dotted green line) than the `STATICMON` baseline (dashed blue line). It is also interesting to notice how the rapid change in the observed trend is not immediately handled by `ADAPTIVEMON`, which shows some delay in sensing the drift and adjusting the sampling rate. In contrast, `STATICMON` always fails to follow the observed time series, confirming the importance of adaptivity in similar contexts.

The *spiky* scenario is the only one resulting in an increment of the RMSE metric for the adaptive Follower (+5.8%). However, this increment is a consequence of the capability to (partially) follow the trend of the indicator. In fact, the `STATICMON` configuration could detect spikes only incidentally, while `ADAPTIVEMON` could change its sampling rate to increase the chance to capture them. Figure 6 illustrates a representative execution of `STATICMON` and `ADAPTIVEMON` for the spiky scenario, with some spikes successfully detected by `ADAPTIVEMON` only. Although successfully capturing some spikes, the reconstructed time series generates a higher error compared to the flat time series reconstructed by `STATICMON`. Indeed, this is a challenging scenario for both approaches (rare short events are hard to detect by monitoring techniques), and more work is required to design cost-effective monitoring techniques that can accurately address spikes.

5.3 RQ2 - Resource Consumption

RQ2 investigates resource consumption considering the two countermeasures currently implemented in `ADAPTIVEMON`, both in isolation and jointly. Again, `STATICMON` is used as the baseline for the comparison. For the experiments reported in this section, we defined a probe (at the follower level only) that exploits the Docker Engine API³ and PowerTOP⁴ to collect the following quality metrics:

- CPU and memory consumption: the percentage of the host’s CPU and memory used.
- NET I/O (MB): the cumulative amount of data sent and received over its network interface from the beginning of the experiment.
- PIDs: the number of processes or threads spawned by the peer.
- PW (mW): the estimated instantaneous power consumption.

The three `ADAPTIVEMON` configurations assessed in this RQ exploit the two strategies defined in Section 3. *Change Rate* adjusts the sampling and forwarding rates of all the collected indicators from 30 to 60 seconds based on the monitored values. *Select Indicators* disables the collection of all indicators except of power consumption if the battery level drops below a threshold. *Combined Countermeasures* uses both strategies. We study the impact of these

Quality Metric	Comparison	Effect Size
<i>CPU consumption</i>	Change Rate vs Combined Countermeasures	small (0.197)
	Select Indicators vs Combined Countermeasures	small (0.260)
<i>Memory consumption</i>	STATICMON vs Select Indicators	large (0.858)
	STATICMON vs Change Rate	large (0.993)
	STATICMON vs Combined Countermeasures	large (0.861)
	Change Rate vs Combined Countermeasures	large (0.990)
<i># spawned sub-processes</i>	Select Indicators vs Combined Countermeasures	small (0.279)
	STATICMON vs Select Indicators	large (0.980)
	STATICMON vs Change Rate	small (0.027)
	STATICMON vs Combined Countermeasures	large (0.963)
<i>Network Input</i>	Change Rate vs Combined Countermeasures	large (0.962)
	STATICMON vs Select Indicators	medium (0.449)
	STATICMON vs Change Rate	small (0.101)
	STATICMON vs Combined Countermeasures	medium (0.483)
<i>Network Output</i>	Change Rate vs Combined Countermeasures	medium (0.431)
	Select Indicators vs Combined Countermeasures	small (0.145)
	STATICMON vs Select Indicators	medium (0.473)
	STATICMON vs Change Rate	small (0.245)
<i>Battery power estimation</i>	STATICMON vs Combined Countermeasures	large (0.521)
	Change Rate vs Combined Countermeasures	medium (0.460)
	Select Indicators vs Combined Countermeasures	small (0.203)
	STATICMON vs Combined Countermeasures	medium (0.376)

Table 3: Statistically valid comparisons for all the quality metrics with their associated effect size.

configurations, along with the `STATICMON` baseline, on resource consumption: a total of four possible configurations are therefore considered. In each experiment, the Follower peer is configured to collect the indicators from its node for 30 minutes and to apply the countermeasures at the beginning of the execution, in such a way the impact of the countermeasures can be accurately measured (in fact we collected more than 300 hundreds samples per metric).

Results. Figure 7 shows a series of five box plots (one for each quality metric) where each plot compares the four compared configurations visually.

We checked the significance of the differences between distributions with the non-parametric Mann–Whitney U test [21], as we observed (via the Shapiro-Wilk test [28]) that such differences are not normally distributed. We specifically checked if the observed differences between the baseline and any other configuration are statistically significant and if *Combined Countermeasures* is significantly better than the individual adaptations strategies (*Change Rate* and *Select Indicators*). We considered a significance level $\alpha = 0.05$, and we also computed the effect size of the observed phenomenon using the Wendt’s formula [39]. Table 3 shows the significant cases only with their corresponding effect size using the conventional categories *small* (less than 0.3), *medium* (between 0.3 and 0.5), and *large* (greater than 0.5).

We only observe marginal differences in CPU consumption. In particular, differences between `ADAPTIVEMON` and `STATICMON` are not significant, and *Combined Countermeasures* introduces significant but small differences compared to employing the other two adaptive strategies individually.

The memory consumption results show statistical significance for all cases with a large effect size for all comparisons, except for *Select Indicators* compared to *Combined Countermeasures* where the effect size is small. The impact of the adaptive strategies on the memory indicator is antithetic: while we notice an increase in

³<https://docs.docker.com/engine/reference/commandline/stats/>

⁴<https://github.com/fenrus75/powertop>




Figure 4: ADAPTIVEMON and STATICMON Follower time series estimations for the *stable-unstable* scenario.




Figure 5: ADAPTIVEMON and STATICMON Leader time series estimations for the *stable-unstable* scenario.




Figure 6: ADAPTIVEMON Follower time series estimation for the *spiky* scenario. The vertical dotted grey lines indicate the *sampling rate*.




Figure 7: STATICMON compared with ADAPTIVEMON countermeasures for each of the collected quality metrics.




Figure 8: STATICMON compared with ADAPTIVEMON countermeasures for the network I/O metrics when the bandwidth is not measured by the Follower.

memory consumption of about 10% for *Change Rate* (compared to STATICMON), memory overhead decreases to about 3% when *Select Indicators* or *Combined Countermeasures* are used. These results can be easily explained considering that although the MAPE-K control loop increases the amount of memory used by ADAPTIVEMON when all probes are active, while when these are disabled (freeing the associated resources) the overall average memory usage decreases.

Since limiting the number of processes can be particularly important when the underlying device platform is resource constrained we measured the number of sub-processes spawned by all compared configurations. In this regard, the number of spawned sub-processes show significant reduction with large effect size when *Select Indicators* and *Combined Countermeasures* are used.

Limiting bandwidth consumption is also extremely important in fog environments, especially in the edge. As a matter of fact, limiting I/O operations is crucial when the network bandwidth is limited and shared by multiple devices and thus can be quickly saturated. Moreover, intensive communication implies high power consumption, a threat to energy efficiency and batteries lifespan in portable devices. The results for network I/O show statistically significant reduction for all adaptive configurations compared to STATICMON, with an effect size ranging from small to medium. Note that results for Input (I) and Output (O) present a similar behavior for the same configuration. More in detail, results show a small effect size for *Change Rate* versus STATICMON comparison. Since we expected a stronger impact of rate adaptation in this context, we analyzed the behavior of the probes, and discovered that the bandwidth gauge exploits iPerf⁵, which measures the bandwidth by saturating it with packets. Such an invasive behavior nullifies the potential benefits of a dynamic sampling rate. Therefore, to further investigate this dimension, we repeated the experiments by disabling the bandwidth monitoring probe for both STATICMON and ADAPTIVEMON. Results are presented in Figure 8. The impact of *Change Rate* is now remarkable, with a reduction on transmitted data ranging between 31% and 34%, with an even higher reduction (between the 37% and the 49%) when both countermeasures are simultaneously active.

⁵<https://iperf.fr/>

Finally, the results on power consumption show meaningful differences only for STATICMON versus *Combined Countermeasures*, suggesting that the individual countermeasures may introduce limited benefits. Still, their combination can significantly improve battery lifetime (with an estimated reduction in power usage of about 36%).

In summary, ADAPTIVEMON can help save network I/O and devices battery without impacting CPU utilization by paying for an extra memory consumption to store the data necessary to run the adaptive mechanisms. This result can be beneficial in the Fog since monitoring systems are called to reduce network overhead without impacting on power consumption [1], especially when devices are battery-powered. At the same time, even resource constrained devices at the edge of the network are usually well equipped in terms of memory, mitigating the impact of a limited extra memory consumed by the monitoring system.

5.4 Threats to Validity

Our study is affected by both internal and external threats to validity. The main internal threats to validity concern with the design of the scenarios used to study RQ1. We proposed five scenarios to mimic different trends. Although indicators collected in real scenarios may behave differently than the ones we investigated, the results obtained with our stereotyped trends are still informative, at least locally (e.g., it is possible to refer to the results reported in the paper for an indicator that becomes unstable or too high).

The definition of the logical states for an indicator depends on several parameters, which are application-dependent. In this paper, we studied how ADAPTIVEMON can be used to obtain self-adaptive capabilities relevant to monitoring, focusing on the assessment of simple countermeasures that do not strongly depend on the domain. Assessing ADAPTIVEMON in the context of dedicated application scenarios is part of our future work and is out of the scope of this paper.

The generality of the results obtained about efficiency (RQ2) might depend on the specific implementation used and the size of the experiment. We used an independent implementation for STATICMON (i.e., FogMon) to minimize any implementation bias, and we added the adaptive behavior to this implementation. To further reduce any implementation threat, we released our solution publicly. In principle, additional experiments may lead to different results. However, we obtained quite clear evidence and we checked the statistical significance of the results to mitigate the risk of overgeneralization.

6 RELATED WORK

In the last decade, a large number of monitoring solutions, both commercial and academic, have been proposed for the Cloud [5, 9, 10, 23, 25, 29, 35, 37, 38]. However, they are seriously challenged by several characteristics of the Fog, such as its massively distributed infrastructure characterized by frequent changes to the topology and the presence of heterogeneous and resource constrained devices [32, 41].

A recent study by Taherizadeh et al. [32] investigated the requirements that must be satisfied by monitoring platforms specialized for adaptive applications orchestrated upon the Cloud-to-Thing

continuum. It is apparent from this survey that none of the solutions currently available for the Cloud can satisfy all the requirements, identifying decentralization and resource optimization via self-adaptation as two of the main open challenges. Similar conclusions have been reported by Abderrahim et al. [1] who explicitly identify the adaptability of the granularity of the reported measurements as one of the key properties for a fog monitoring system.

Monitoring approaches specifically designed for the fog environment have been recently investigated [8, 13, 15, 30]. In particular, FMonE [8] is a monitoring tool that relies on a container orchestration system to build monitoring pipelines, addressing the distinctive features of a fog infrastructure. It provides users with the flexibility to define their monitoring pipelines and operate them across the active regions. The container orchestration system is the (centralized) module of the FMonE architecture responsible for launching the monitoring agents and the data backends in each region.

PyMon [15] is a lightweight prototypical monitoring framework available for relevant Docker-enabled architectures such as ARM, AARCH64 and x86_64, and particularly suitable for single board computers (SBC) at the edge of the network. It extends the host-based monitoring tool Monit with capabilities to inspect running Docker containers.

Souza et al. [30] proposed a monitoring tool that extends the CLABS model [3] and is capable of monitoring resources by deploying services between the Edge and the Cloud.

Unlike ADAPTIVEMON, none of these solutions implement adaptive policies to adapt the behavior of the monitored system to the collected data. Furthermore, they are not based on a P2P architecture, so that, they struggle to cope with some of the fog distinctive traits such as heavily distributed infrastructures, rapid changes in the topology, and communication links failures.

FogMon [13] is the fog-oriented monitoring framework that we extended to implement ADAPTIVEMON. It collects and aggregates data about resource consumption, network conditions, and IoT devices connected to fog nodes. It exploits a two-tier (Leader-Follower) P2P architecture and gossip protocols to reduce the network overhead. Also, it adapts the number of Leader nodes in the P2P overlay and the underlying Followers topology based on current network conditions. However, it does not provide self-adaptive behaviors to govern the internal functioning of the monitoring system as designed in ADAPTIVEMON. For instance, FogMon cannot be used to dynamically change the set of the collected indicators or the sampling rate.

Among the works that are not specifically designed for the Fog, it is worth mentioning some that still relate to ADAPTIVEMON. In particular, ADMIn [36] is an IoT-specific monitoring framework designed to reduce devices' energy consumption and the volume of data sent over the network. This is achieved essentially by adapting the rate at which devices disseminate monitoring streams based on run-time knowledge (e.g., stream evolution, variability, seasonality).

Also, Tangari et al. [33] propose a self-adaptive and decentralized framework for resource monitoring in the scope of Software Defined Networks (SDN). It enables metrics collection through a self-tuning and adaptive monitoring technique that adjusts its settings based on traffic dynamics to balance operation costs with monitoring accuracy while reducing network overhead. Compared to ADAPTIVEMON, the proposed framework lacks generality since

the adaptation capabilities are limited to some predetermined aspects and is not designed to support the capability to run multiple and diverse adaptation rules.

Finally, SkyEye [14] is a tree-based monitoring solution operating on structured P2P overlay networks. It provides continuous monitoring for a wide range of metrics for all peers in the network. It is characterized by a tree structure, which enables peer partitions in a hierarchical fashion. The aggregated statistics received by the upper layers of the tree describe the statistics of the peers in the corresponding sub-trees. Messages are used to disseminate the global statistics retrieved from the top levels and maintain the tree topology. However, unlike ADAPTIVEMON it is not explicitly designed for the Fog, and it completely lacks adaptivity.

7 CONCLUSIONS

Fog monitoring calls for approaches that can deal with the dynamism of the environment while taking the available resources into consideration. P2P monitoring approaches cope with some of the architectural-level traits of the Fog, such as nodes that can dynamically join and leave the network or malfunctioning in their communication links, but do not provide self-adaptive capabilities that can be used to dynamically adapt the behavior of the monitoring system to changing conditions.

In this paper, we presented ADAPTIVEMON, a self-adaptive P2P monitoring solution that exploits a knowledge base continuously fed with monitoring data to dynamically change the system behavior by activating countermeasures. Experimental results show that adaptive behaviors can improve monitoring accuracy while optimizing the usage of the available resources, in comparison to a non-adaptive solution.

Future work will focus on extending the self-adaptive capabilities of ADAPTIVEMON by considering several application scenarios, including anomaly detection and self-healing, and multiple application domains such as Service Level Agreements (SLAs) and Function-as-a-Service (FaaS) monitoring in the fog environment. We also plan to study how to exploit the MAPE-K paradigm at the level of the whole P2P monitoring system, to deliver adaptive capabilities at the system level (e.g., to run adaptation procedures that require the synchronized intervention of multiple peers), instead of the individual peers only.

REFERENCES

- [1] Mohamed Abderrahim, Meryem Ouzzif, Karine Guilloaud, Jerome Francois, and Adrien Lebre. 2017. A holistic monitoring service for fog/edge infrastructures: a foresight study. In *2017 IEEE 5th International Conference on Future Internet of Things and Cloud (FiCloud)*. IEEE, 337–344.
- [2] Giuseppe Aceto, Alessio Botta, Walter De Donato, and Antonio Pescapé. 2013. Cloud monitoring: A survey. *Computer Networks* 57, 9 (2013), 2093–2115.
- [3] Khalid Alhamzani, Rajiv Ranjan, Prem Prakash Jayaraman, Karan Mitra, Chang Liu, Fethi Rabhi, Dimitrios Georgakopoulos, and Lizhe Wang. 2015. Cross-layer multi-cloud real-time application QoS monitoring and benchmarking as-a-service framework. *IEEE Transactions on Cloud Computing* 7, 1 (2015), 48–61.
- [4] Mohammed Alodib. 2016. QoS-Aware approach to monitor violations of SLAs in the IoT. *Journal of Innovation in Digital Ecosystems* 3, 2 (2016), 197–207.
- [5] Wolfgang Barth. 2008. *Nagios: System and network monitoring*. No Starch Press.
- [6] Theophilus Benson, Sambit Sahu, Aditya Akella, and Anees Shaikh. 2010. A First Look at Problems in the Cloud. *HotCloud* 10 (2010), 15.
- [7] 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*. 13–16.
- [8] Álvaro Brandón, María S Pérez, Jesus Montes, and Alberto Sanchez. 2018. Fmone: A flexible monitoring solution at the edge. *Wireless Communications and Mobile Computing* 2018 (2018).
- [9] Elasticsearch B.V. 2022. Elastic Stack: Elasticsearch, Kibana, Beats & Logstash. <https://www.elastic.co/elastic-stack/> Accessed on 20th of January 2022.
- [10] Jose M Alcaraz Calero and Juan Gutierrez Aguado. 2014. Monpaas: An adaptive monitoring platform as a service for cloud computing infrastructures and services. *IEEE Transactions on Services Computing* 8, 1 (2014), 65–78.
- [11] Michele Ciavotta, Davide Motterlini, Marco Savi, and Alessandro Tundo. 2021. DFaaS: Decentralized Function-as-a-Service for Federated Edge Computing. In *2021 IEEE 10th International Conference on Cloud Networking (CloudNet)*, 1–4. <https://doi.org/10.1109/CloudNet53349.2021.9657141>
- [12] Dalibor Dobrilovic, Francesco Flammini, Andrea Gaglione, and Daniel Tokody. 2021. Open-source hardware performance in Edge Computing. <https://smartcities.ieee.org/newsletter/july-2021/open-source-hardware-performance-in-edge-computing> Accessed on 20th of January 2022.
- [13] Stefano Forti, Marco Gaglianese, and Antonio Brogi. 2021. Lightweight self-organising distributed monitoring of Fog infrastructures. *Future Generation Computer Systems* 114 (2021), 605–618.
- [14] Kalman Graffi and Andreas Disterhöft. 2017. SkyEye: A tree-based peer-to-peer monitoring approach. *Pervasive and Mobile Computing* 40 (2017), 593–610.
- [15] Marcel Großmann and Clemens Klug. 2017. Monitoring container services at the network edge. In *2017 29th International Teletraffic Congress (ITC 29)*, Vol. 1. IEEE, 130–133.
- [16] Pengfei Hu, Sahraoui Dhelim, Huansheng Ning, and Tie Qiu. 2017. Survey on fog computing: architecture, key technologies, applications and open issues. *Journal of network and computer applications* 98 (2017), 27–42.
- [17] Márk Jelasity. 2011. Gossip. In *Self-organising software*. Springer, 139–162.
- [18] Steven J Johnston, Philip J Basford, Colin S Perkins, Herry Herry, Fung Po Tso, Dimitrios Pazaros, Robert D Mullins, Eiko Yoneki, Simon J Cox, and Jeremy Singer. 2018. Commodity single board computer clusters and their applications. *Future Generation Computer Systems* 89 (2018), 201–212.
- [19] Jeffrey O Kephart and David M Chess. 2003. The vision of autonomic computing. *Computer* 36, 1 (2003), 41–50.
- [20] Michael Kleis, Eng Keong Lua, and Xiaoming Zhou. 2005. Hierarchical peer-to-peer networks using lightweight superpeer topologies. In *10th IEEE Symposium on Computers and Communications (ISCC'05)*. IEEE, 143–148.
- [21] Henry B Mann and Donald R Whitney. 1947. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics* (1947), 50–60.
- [22] Dejan S Milojicic, Vana Kalogeraki, Rajan Lukose, Kiran Nagaraja, Jim Pruyne, Bruno Richard, Sami Rollins, and Zhichen Xu. 2002. Peer-to-peer computing.
- [23] Rihards Olups. 2016. *Zabbix Network Monitoring*. Packt Publishing Ltd.
- [24] Andy Oram. 2001. *Peer-to-Peer: Harnessing the power of disruptive technologies*. " O'Reilly Media, Inc."
- [25] Javier Povedano-Molina, Jose M Lopez-Vega, Juan M Lopez-Soler, Antonio Corradi, and Luca Foschini. 2013. DARGOS: A highly adaptable and scalable monitoring architecture for multi-tenant Clouds. *Future Generation Computer Systems* 29, 8 (2013), 2041–2056.
- [26] Gary Riley. 1991. Clips: An expert system building tool. In *Proceedings Second National Technology Transfer Conference and Exposition*, Vol. 2. 149–158.
- [27] Seedstudio.com. 2020. All about cpus: Microprocessor, microcontroller and Single Board Computer. <https://www.seedstudio.com/blog/2020/10/27/all-about-cpus-microprocessor-microcontroller-and-single-board-computer/> Accessed on 20th of January 2022.
- [28] Samuel Sanford Shapiro and Martin B Wilk. 1965. An analysis of variance test for normality (complete samples). *Biometrika* 52, 3/4 (1965), 591–611.
- [29] Anas Shatnawi, Matteo Orrù, Marco Mobilio, Oliviero Riganelli, and Leonardo Mariani. 2018. Cloudhealth: a model-driven approach to watch the health of cloud services. In *2018 IEEE/ACM 1st International Workshop on Software Health (SoHeal)*. IEEE, 40–47.
- [30] Arthur Souza, Nélio Cacho, Ayman Noor, Prem Prakash Jayaraman, Alexander Romanovsky, and Rajiv Ranjan. 2018. Osmotic monitoring of microservices between the edge and cloud. In *2018 IEEE 20th International Conference on High Performance Computing and Communications; IEEE 16th International Conference on Smart City; IEEE 4th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*. IEEE, 758–765.
- [31] Ralf Steinmetz and Klaus Wehrle. 2005. *Peer-to-peer systems and applications*. Vol. 3485. Springer.
- [32] Salman Taherizadeh, Andrew C Jones, Ian Taylor, Zhiming Zhao, and Vlado Stankovski. 2018. Monitoring self-adaptive applications within edge computing frameworks: A state-of-the-art review. *Journal of Systems and Software* 136 (2018), 19–38.
- [33] Gioacchino Tangari, Daphne Tuncer, Marinos Charalambides, Yuanshunle Qi, and George Pavlou. 2018. Self-adaptive decentralized monitoring in software-defined networks. *IEEE Transactions on Network and Service Management* 15, 4 (2018), 1277–1291.
- [34] David Tracey and Cormac Sreenan. 2019. How to see through the Fog? Using Peer to Peer (P2P) for the Internet of Things. In *2019 IEEE 5th World Forum on Internet of Things (WF-IoT)*. IEEE, 47–52.
- [35] Demetris Trihinas, George Pallis, and Marios D Dikaiakos. 2014. Jcatascopia: Monitoring elastically adaptive applications in the cloud. In *2014 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*. IEEE, 226–235.
- [36] Demetris Trihinas, George Pallis, and Marios D Dikaiakos. 2017. ADMIN: Adaptive monitoring dissemination for the Internet of Things. In *IEEE INFOCOM 2017-IEEE conference on computer communications*. IEEE, 1–9.
- [37] Alessandro Tundo, Marco Mobilio, Matteo Orrù, Oliviero Riganelli, Michell Guzmàn, and Leonardo Mariani. 2019. Varys: An agnostic model-driven monitoring-as-a-service framework for the cloud. In *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1085–1089.
- [38] James Turnbull. 2018. *Monitoring with Prometheus*. Turnbull Press.
- [39] Hans W Wendt. 1972. Dealing with a common problem in social science: A simplified rank-biserial coefficient of correlation based on the statistic. *European Journal of Social Psychology* (1972).
- [40] B Beverly Yang and Hector Garcia-Molina. 2003. Designing a super-peer network. In *Proceedings 19th international conference on data engineering (Cat. No. 03CH37405)*. IEEE, 49–60.
- [41] Ashkan Yousefpour, Caleb Fung, Tam Nguyen, Krishna Kadiyala, Fatemeh Jalali, Amirreza Niakanlahiji, Jian Kong, and Jason P Jue. 2019. All one needs to know about fog computing and related edge computing paradigms: A complete survey. *Journal of Systems Architecture* 98 (2019), 289–330.