

Regular Paper: AQuA-CEP: Adaptive Quality-Aware Complex Event Processing in the Internet of Things

ABSTRACT

Sensory data has a profound influence on the quality of detected events in a distributed complex event processing system (DCEP). Since each sensor's status is not stable at runtime, a single sensing assignment is often not enough to fulfill the consumer's quality requirements. In this paper, we study in the context of AQuA-CEP the problem of dynamic quality monitoring and adaptation of suitable data sources. To support this, in AQuA-CEP, queries to detect complex events are supplemented with consumer-definable quality policies that are evaluated and used to autonomously select (or even configure) suitable data sources of the sensing infrastructure. In addition, we studied different forms of expressing quality policies and analyzed how it affects the quality monitoring process. Various modes of evaluating and applying quality-related adaptations and their impacts on correlation efficiency are addressed, too. We assessed the performance of AQuA-CEP in IoT scenarios by utilizing the notion of the quality policy alongside the query processing adaptation using knowledge derived from quality monitoring. The results show that AQuA-CEP can improve the performance of DCEP systems in terms of the quality of results while fulfilling the consumer's quality requirements. Quality-based adaptation can also increase the network's lifetime by optimizing the sensor's energy consumption due to efficient data source selection.

CCS Concepts: • Computer systems organization → Real-time systems; • Networks → Network dynamics; • Software and its engineering → Publish-subscribe / event-based architectures.

Keywords: Complex Event Processing, Stream Processing, Adaptation, Quality, Internet of Things.

ACM Reference Format:

. 2022. Regular Paper: AQuA-CEP: Adaptive Quality-Aware Complex Event Processing in the Internet of Things. In *Proceedings of ACM Conference (DEBS'17)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DEBS'17, June 27-30, 2023, Neuchatel, Switzerland

© 2022 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM. . \$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Reacting to varying situations is a fundamental requirement in the Internet of Things (IoT) scenarios like traffic monitoring, healthcare systems, and smart homes. Distributed Complex Event Processing (DCEP) is a widely employed paradigm to support efficient situation detection based on a variety of distinct sensors and a step-wise transformation from *primary events* to situations of interest for consumers in the form of *complex events*. The resulting quality usually expressed in the form of Quality of Service (QoS) and Quality of Results (QoR), highly depends on the origin of primary events. Especially in IoT scenarios in which the primary events are generated often based on distributed sensor readings from the environment. These sensing deployments are vulnerable to the immense dynamicity that existed in the environment (e.g., availability of the sensors) and a single sensing deployment is often not enough to meet quality requirements related to the system and its consumers.

An established solution to react properly to environmental dynamics is to adapt the detection logic's placement to the available computing resources which are part of the DCEP framework. Also, such an adaptation needs to deal with the limitations that the allocated resources might have during the query execution (cf. [6, 18–20, 29]). This idea provides essential means to maintain or improve the QoS-related measures, e.g., by reducing the imposed end-to-end delay or regulating the bandwidth consumption. On the other hand, the idea of sacrificing QoR to keep QoS at an acceptable level can already benefit the DCEP systems by combining these mechanisms with other runtime approaches like load shedding techniques [26, 27]. Although influencing QoR will lead to a degradation in consumer requirements' satisfaction, these techniques find it crucial to impact QoR as less as possible, e.g., by dropping events from partial matches.

The state-of-the-art approaches decouple the detection procedure and adaptation strategies from sensing deployment configuration and operate only based on information existing in the design time. Such an idea limits the system's capabilities to react properly to dynamics in the sensing layer, e.g., the quality of sensor readings or their battery level. Hence, the degradation in QoR can be propagated due to these limitations while the DCEP system is not able to actively influence and prevent the consequences of hardwired sensing configuration. Moreover, even if the sensing deployment adaptation is considered at runtime, adaptation strategies' outcome can affect the QoR [12]. For example, in the case of updating the data sources from a camera to a motion detector, the motion event's accuracy would be degraded. In this vein, due to attaining the most elevated quality grade in the adaptation decisions, the inputs are mandated to have an admissible level of quality. In this regard, the input data can be

assumed of insufficient quality if not accurate, precise, fresh, or truthful. Events are also evaluated as inadequate quality if they do not hold a certain level of confidence, are received out of order, are wrongly detected, or are not detected at all. Therefore, measuring to what extent the consumer's quality requirements are met should be taken into account when applying any adaptation strategy.

In this paper, we analyze and present concepts on expanding flexibility and adaptivity by proposing quality-aware event processing to enhance QoR and QoS. In particular, AQuA-CEP is a mechanism that is designed based on the idea of dynamically exchanging sensing deployment concerning the demarcated requirements by the consumer that influence how sensor data is processed. We enhance DCEP with the concept of so-called quality policies and corresponding quality monitoring mechanisms. Upon any change occurring in the environment or in the sensing deployment observed by our DCEP system, AQuA-CEP will autonomously adapt the sensing deployment according to the available sensing infrastructure, if required. It can be performed by defining sensing configuration restrictions (e.g., cost constraints) considering quality requirements expressed by consumers. Consequently, an efficient sensing deployment assignment is performed by a utility metric concerning the defined restrictions. For more details, AQuA-CEP provides the following contributions:

1. We provide a new representation of the quality demands of a query in DCEP systems by proposing a policy-driven specification of complex events to boost data processing performance and more promising utilization of IoT resources.
2. We devise how quality monitoring can be applied in DCEP by presenting concepts allowing the dynamic reconfiguration of appropriate data sources while fulfilling consumer's quality requirements.
3. We explore strategies for configuring quality monitoring agents that triggers adaptation strategies upon any quality policy violation and address the impacts each configuration might have on the DCEP system's performance in terms of QoS and QoR.
4. We evaluate the performance of our proposed mechanism with a real-world dataset alongside using synthetic data to show the ability of AQuA-CEP to boost the performance of the DCEP system in adapting the sensing deployment while observing consumer quality requirements.

The remainder of this paper is structured as follows. We introduce an overview of the AQuA-CEP system model in Section 2. We elaborate on the problem statement in Section 3. In Section 4, the detailed overview of AQuA-CEP is presented. The evaluation results of AQuA-CEP are exhibited in Section 5. The related work is presented in Section 6. Finally, Section 7 concludes our paper and points to our future work.

2 SYSTEM MODEL

In this section, we present the system model by introducing AQuA-CEP components, the model of IoT devices, and the provided adaptation models.

2.1 AQuA-CEP Model

We consider a DCEP system to consist of multiple *producers* (e.g., mobile phones, etc.) that generate streams of primary events from the received sensory data and announce them to the system as their advertisements. Correspondingly, *consumers* (e.g., users, applications, services, etc.) create situations of interest as subscriptions and submit them to the system as continuous queries where the set $Q = \{q_1, \dots, q_n\}$ denotes the set of currently deployed queries. Moreover, a group of *brokers* (i.e., CEP engines) performs computational tasks (e.g., filter, join, etc.) by hosting a set of *operators* and forward the result to the next step that can be another operator or the consumers.

A query q_i explains the logic by which a complex event can be detected over primary event streams. It can be performed by applying standard CEP operators like pattern matching, aggregation, or windowing over primary events or their attributes. To do so, the imposed complex event detection logic should be applied to the specific brokers for execution. In the meantime, consumers are also allowed to specify their quality requirements as part of the query (e.g., the location accuracy of less than one meter). We denote the set of *consumer-side constraints* of all deployed queries in Q by $\mathcal{G} = \{g_1, g_2, \dots, g_k\}$. A data source is considered an eligible candidate to feed the system if it can meet all related consumer-side constraints.

2.2 IoT Resource Model

In AQuA-CEP, sensors are the origin of data that measure a specific phenomenon (e.g., temperature) in the environment. The sensory data sources (e.g., Bluetooth) that are used at a given time t form the set of active sensing deployments $\mathcal{SD} = \{sd_1, \dots, sd_j\}$. Here, sd_i refers to a specific data source among all available options in the environment that can be participated reliably in the process of sensor assignment for a deployed query. The availability of data sources is dynamic meaning the set \mathcal{SD} might change over time. In the IoT environment, a data source can be mobile (e.g., sensors embedded in a smartphone) or stationary (e.g., a surveillance camera). We assume that the mobility status of data sources does not negatively or positively affect the quality of their readings.

In our scenario, IoT devices (e.g., smartphones) are interconnected to the system over a wireless sensor network and demanded to register their sensing deployment in the AQuA-CEP in advance. These devices represent the CEP producers who generate primary event streams from sensory data. Also, some of these devices are eligible to issue queries acting as CEP consumers. Moreover, CEP operators can be placed on IoT resources with sufficient computing capabilities, e.g., on the cloud or fog nodes.

2.3 Adaptation Model

For coordinating adaptation and selecting its correct triggers, AQuA-CEP needs to monitor the quality of produced events as well as the state of the sensing infrastructure. In this regard, we build on a sensing middleware (e.g., [1]) that offers the possibility to identify, configure, and access the physical sensors. The system samples the quality level

for a subset of the produced events and evaluates potential alternative configurations.

In AQuA-CEP, the adaptation decisions need to serve multiple objectives, e.g., the cost for fulfilling the query's quality policies includes the expenses for utilizing the sensing infrastructure and completing reconfigurations. Besides, adaptation should guarantee a level of stability which means how often the achieved quality of the detected event stays inside a predefined threshold region after applying the adaptation strategy. It would avoid oscillation and inessential switching costs. Also, with AQuA-CEP we are looking to adapt the event processing to the environmental dynamics by switching the sensing deployment. For example, given a new sensor registered in the network by which some of the currently running queries would be answered. In this case, AQuA-CEP generates new query models considering the new sensor and checks for the costs imposed by a transition from the current sensing deployment to the new one.

3 PROBLEM STATEMENT

In this work, AQuA-CEP as a DCEP system selects suitable data sources, i.e., a sensing deployment $\alpha(\mathcal{SD}) \subset \mathcal{SD}$, where α determines which sensor sources of \mathcal{SD} will be used. AQuA-CEP is required to meet the consumer constraints in quality or notify consumers when no proper sensing deployment is feasible. Furthermore, each sensor source sd_i of a sensing deployment imposes a *System-Side Cost (SSC)* denoted by $C_{SSC}(sd_i)$ as well as the cost for performing *Quality Monitoring (QM)* for every query q_k denoted by $C_{QM}(sd_i, q_k)$. C_{SSC} includes those quality metrics that are more important for the system (i.e., energy consumption, reusability, resource utilization, etc.). C_{QM} is the cost imposed by the configuration model of the monitoring agent in terms of time (i.e., the delay related to the processing events in the operator and delay for performing the transition between sensing deployments) and computation (i.e., the required computation resources to conduct monitoring process).

More formally, AQuA-CEP aims to find α which minimizes the cost factors imposed by system-side costs and quality monitoring costs subject to the quality constraints of a consumer, i.e.,

$$\begin{aligned}
 \min \quad & w_s \sum_{sd_i \in \mathcal{SD}} \alpha(sd_i) C_{SSC}(sd_i) \\
 & + w_q \sum_{sd_i \in \mathcal{SD}} \alpha(sd_i) \sum_{q_k \in \mathcal{Q}} C_{QM}(sd_i, q_k) \\
 \text{s.t.} \quad & \alpha(\mathcal{SD}) \text{ satisfies constraints in } \mathcal{G} \\
 & \alpha(sd_i) = 1 \text{ iff } sd_i \text{ is selected.} \\
 & \alpha(sd_i) \in \{0, 1\}
 \end{aligned}$$

where w_s indicates the weight related to system-side costs, and w_q is the weight associated with monitoring costs.

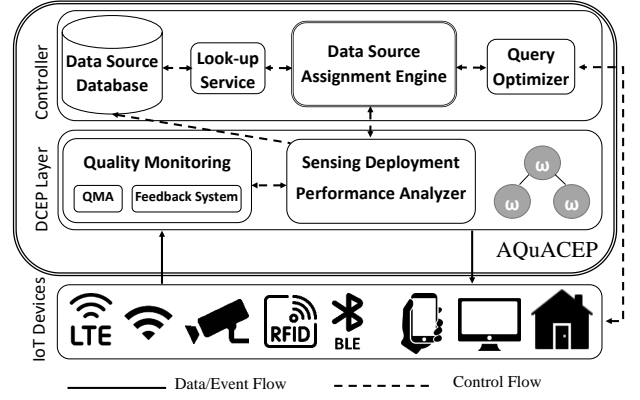


Figure 1. The AQuA-CEP System Design.

4 THE AQuA-CEP SYSTEM DESIGN

In Figure 1, we depict the foundational components employed in AQuA-CEP. By utilizing Software-Defined Networking (SDN), we make use of a *controller* to function as the coordinator module and enforce a quality-driven DCEP. This component is logically centralized but physically distributed and is in charge of exchanging control messages to synchronize the event detection procedure. To do so, the controller owns the principal role in matching the subscriptions to advertisements. A *query optimizer* component receives queries submitted by consumers, transforms the query description to a set of event types, and passes the list of required data sources to the *data source assignment engine*. A *look-up service* is triggered by the engine to explore the potential candidates for each event type in *data source database*, where the currently available data sources are previously registered themselves. The records in this database are dynamic and can be registered or canceled at runtime.

Moreover, the controller's functionality is enriched by employing *quality monitoring* agents in the DCEP layer. Upon any predefined situation (e.g., data source disconnection), a so-called *control event* is created by the responsible agent. It notifies the controller to execute corresponding steps as adaptation scenarios in order to maintain the QoR. To do so, a *sensing deployment performance analyzer* investigates the current state of assigned data sources using information acquired by quality monitoring agents and updates the assignment engine to reconfigure sensing deployment, if necessary. The performance analyzer component also updates the data source database's records based on the quality monitoring results. It influences the characteristics of data sources or their availability.

Besides, parts of our system are built on existing concepts for the flexible execution of event processing operators, as proposed in TCEP [19] and CEPLESS [17]. This allows AQuA-CEP to modify the deployment and configuration of operators and integrate a wide range of additional event processing engines, e.g., Apache Flink. With such flexibility, AQuA-CEP can revise the operators' deployment and influence important QoS characteristics, e.g., bandwidth usage.

4.1 Quality Requirement Description

The foremost step in acquiring the consumers' quality expectations is the technique by which they can elaborate on their requirements. Such a method must be not only easy to use for the consumers, but also sufficiently comprehensive to cover all aspects and flexible to fit different types of consumers' quality requirements. Hence, we enable consumers to express their specific requirements related to a given query in form of *quality policy*.

Quality Policy (QP). AQuA-CEP extends traditional query specification of event processing systems by providing the possibility to specify quality requirements. A quality policy can determine diverse quality metrics which are essential for the consumers, e.g., the accuracy of delivered values as part of the event detection process. Also, a quality policy might comprise *threshold levels* and *priorities* that can be exploited to optimize and trade between conflicting demands of multiple deployed queries. According to the type of thresholds, the quality policy can be categorized as *static* or *dynamic*.

Quality requirements in the static type of quality policies are specified based on static thresholds as the exact amounts stated clearly in the query, e.g., the temperature data is requested to be delivered with an interval of 10 seconds. Only one type of quality metric can be involved in each static quality policy. Therefore, to assess multiple aspects of each event one expression is required in the query definition. For each data type, the system needs to provide a manner for the consumers to define acceptable values. For instance, to determine the resolution of images, the consumer should have the possibility to specify the thresholds based on PPI (i.e., pixels per inch). Thus, the expression of quality policies for consumers will be as easy as possible.

Some sorts of quality requirements cannot be represented by a static threshold. In those circumstances, we define a dynamic threshold that varies depending on a second factor which can be time or a context-related parameter. For instance, a dynamic threshold based on the location factor looks like "the location accuracy of an object should be less than 2 meters if it is within 100 meters of a particular area. Otherwise, 10-meter accuracy would be sufficient". By utilizing dynamic thresholds, more intricate descriptions for quality requirements can be explainable enhancing the flexibility of query definition. Such sort of flexibility will improve consumer satisfaction while optimizing our data analytic system to avoid utilizing more complex procedures to fulfill quality requirements.

In order to validate the admissibility of thresholds, the controller inspects the capability of the data sources and adjusts their characteristics based on the requested thresholds in the query, if applicable. In case there are no appropriate data sources available in the environment concerning the quality requirements, the controller revises the query model with the acceptable thresholds and notifies the consumer about the new query model. Then, based on the feedback obtained from the consumer, the controller will deploy the newly produced query model or cancel the query processing. On the other hand, priorities can also

be determined in the query definition. It is worth noticing that a higher query priority will impose higher costs for the query issuer. Such costs can be defined by the system designer depending on the use case.

Moreover, the data source conditions and consumer's quality requirements might be varied over time. That is why the quality policies have to be revised during the runtime. For example, the consumer may need the result of a running query very urgently, such as the current blood pressure of a patient who may have an acute condition with lower intervals. Therefore, the consumer needs to inform the system about this change by raising the query's priority as well as changing the sensing interval's threshold. To do so, AQuA-CEP employs a feedback process for maintaining and updating the quality policies and renewing the policies whenever the quality requirements or the data sources' status is altered.

4.2 Quality Monitoring

Monitoring the quality of produced results is essential for AQuA-CEP to ensure the expressed quality requirements can be met and adaptation decisions are conducted timely. In stream processing, quality observation should be performed with the lowest possible delay. On the other hand, the available resources for computation are usually bounded. Thus, an event monitoring process must ponder both of these aspects simultaneously.

4.2.1 Quality Management Agent (QMA). One of the novel traits of AQuA-CEP is to employ QMAs that are liable for inspecting the predefined quality metrics over the event streams and triggering warnings in the matter of quality degradation. On the other hand, the utilization of QMA and where it is hosted in the data plane might influence the quality of service. So, we defined the concept of QMA's configuration model and discussed how it improves the quality in various processing stages.

4.2.2 QMA Configuration Models. As we discussed earlier, the form of QMA configuration in the DCEP layer can yield different results. There are two kinds of configuration models; *sequential* and *parallel* as shown in Figure 2. In a sequential configuration model, the events from all producers that fulfill the requirements of a specific query are aggregated into one joint event stream and fed into the corresponding QMA for quality evaluation. In this case, QMA is in charge of filtering events and allows those events that possess the required level of quality delivered to the respective CEP engine. The primary advantage of this model is that the transition time (i.e., handover), that is the time required to swap data sources, is ensured to remain small, i.e., in the order of milliseconds. Moreover, it is feasible joining events from two or more producers to boost the quality by utilizing redundancy. On the other hand, processing all events from the producers certainly imposes considerable latency in query processing that should be taken into account.

The parallel configuration model will impose the minimum possible latency in processing since QMA analyzes the event attributes' quality in parallel. However, it is only possible to connect one producer to the respective

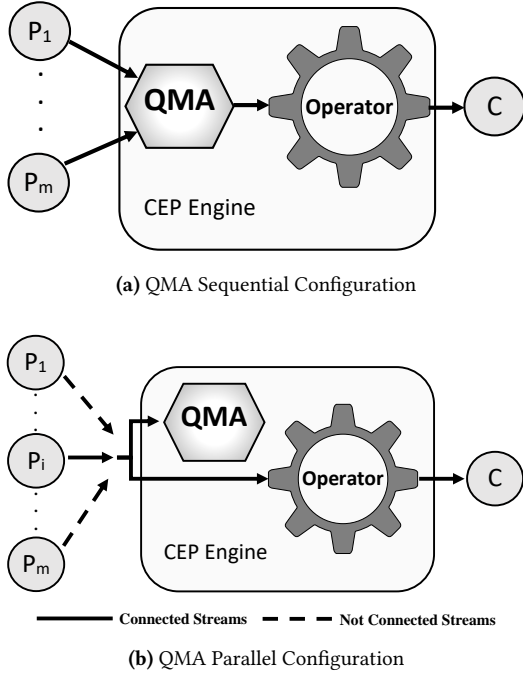


Figure 2. QMA Configuration Models.

operator. Thus, if the quality of the event produced by this data source degrades, it takes time for the QMA to trigger an alert and request the controller to link another data source to the CEP engine that fulfills the quality requirements. Such a transition surely imposes a noticeable overhead on the event processing system. Nevertheless, both processing and transition delay should be considered as the major costs when the system wants to decide on the configuration model of QMAs.

Moreover, the QMA costs for monitoring each data source can be reviewed based on two factors; time and computation, directly dependent on the QMA's configuration model. In terms of time, if the parallel configuration model is chosen, the cost is the delay related to switching between the current data source (i.e., sd_i) and the next option (i.e., sd_l) as $C_s(sd_i, sd_l, q_k)$. On the contrary, if the sequential configuration model was chosen, the time overhead is the delay caused by quality analysis ($C_a(sd_i, q_k)$). In terms of computation costs, the overhead in both parallel and sequential models is almost the same. We called this ($C_w(sd_i, q_k)$) includes the required resources for analyzing an event stream using sliding windows and the computing resources for the data source assignment ($C_{ads}(q_k)$). Therefore, the total cost of monitoring the data source sd_i using a QMA is as follows.

$$\begin{aligned}
 C_{QM}(sd_i, q_k) = & w_{seq} C_s(sd_i, sd_l, q_k) \\
 & + w_{par} C_a(sd_i, q_k) \\
 & + C_w(sd_i, q_k) + C_{ads}(q_k)
 \end{aligned}$$

$$\begin{aligned}
 s.t. \quad w_{seq} = 1 & \text{ iff 'Sequential mode is selected.'} \\
 w_{par} = 1 & \text{ iff 'Parallel mode is selected.'} \\
 w_{seq} \in \{0, 1\} & \text{ and } w_{par} \in \{0, 1\}
 \end{aligned}$$

4.3 Sensing Deployment Adaptation

Adaptation decisions in AQuA-CEP are performed following the MAPE-K feedback loop model [4] building on the previous two steps (i.e., quality description and monitoring). Consequently, the monitoring outcomes are analyzed within every loop to determine adaptation decisions and finally apply them to the stream processing infrastructure.

4.3.1 Controller. In Algorithm 1, we represent the controller's core functionality in AQuA-CEP.

Algorithm 1 Controller Functionality

```

1: Initialization:
    $Q_{Qid} = et_1, \dots, et_l \leftarrow$  User Query Qid;
    $QP \leftarrow \{qp_1, \dots, qp_n\}$ ;
    $SUB \leftarrow \emptyset$ ;
    $ADV \leftarrow \emptyset$ ;
2: upon ( $C_{cid}$ .Submit( $Q_{Qid}$ )) do
3:   for  $et_i \in Q_{Qid}$  do
4:      $SUB \leftarrow SUB \cup sub_{Qid}^{et_i}$ ;
5:   AssignDataSource( $ADV$ ,  $SUB$ );  $\triangleright$  Algorithm 2
6:   upon ( $P_{pid}$ .Register( $\mathcal{SD}(P_{pid})$ )) do
7:     for  $sd_i \in \mathcal{SD}(P_{pid})$  do
8:        $ADV \leftarrow ADV \cup adv_{pid}^{sd_i}$ ;
9:     AssignDataSource( $ADV$ ,  $SUB$ );  $\triangleright$  Algorithm 2
10:  upon (QMA.Alarm) do
11:    ProcessAlarm(QMA.Alarm);  $\triangleright$  Algorithm 3
    
```

The first step of query processing is to initialize the corresponding variables, i.e., the set of event types, their quality policies, and the thresholds for each of these policies that need to be clearly identified by the consumer. When the consumer with ID C_{cid} registers the query, the controller initiates a subscription for each query's simple event. Such a subscription comprises details regarding the event type of simple event (i.e., et_i) and the respective query (i.e., Qid). On the other hand, a producer P_{pid} registers its available sensing deployments (i.e., $\mathcal{SD}(P_{pid})$) in the controller. When the producer completes the registration procedure, the controller generates an advertisement for each sensing deployment of the producer. It includes information about the respective data source and the ID of the producer. After submitting each query or registering data sources, the controller investigates possible solutions to match current advertisements to subscriptions (Refer to Algorithm 2).

Moreover, AQuA-CEP aims to assign data sources adaptively according to 1) changes in the environment, 2) quality of data streams and 3) dynamic thresholds of quality policies. Hence, a feedback system continuously inspects the conditions in both the data plane and the control plane to trigger alarms upon any noteworthy changes. Upon accepting any warning from QMAs which indicates a quality-related situation, AQuA-CEP examines the alarm type and processes it promptly. Depending on the type, AQuA-CEP decides which reconfiguration should apply in the data source assignment (Refer to Algorithm 3).

4.3.2 Data Source Management. In Algorithm 2, we represent the process of assigning data sources.

Algorithm 2 Data Source Management

```

1: function ASSIGNDATASOURCE( $ADV, SUB$ )
2:   for  $sub^{et_i} \in SUB$  do
3:      $QP_{sub^{et_i}} \leftarrow$  Related quality policies to  $et_i$ ;
4:      $M_{ADV^{et_i}} \leftarrow$  Matching advertisements to  $et_i$ ;
5:     for  $adv^{sd_i} \in M_{ADV^{et_i}}$  do
6:       if  $adv^{sd_i}.MeetAllQP(QP_{sub^{et_i}})$  then
7:          $C_{sd_i} \leftarrow C_{SSC}(sd_i) + C_{QMA}(sd_i, q_k)$ ;
8:          $\mathcal{SD}(sub^{et_i}) \leftarrow \mathcal{SD}(sub^{et_i}) \cup (adv^{sd_i},$ 
            $C_{sd_i})$ ;
9:    $\alpha_t(\mathcal{SD}) \leftarrow$  HACS( $\mathcal{SD}, \mathcal{G}$ );            $\triangleright$  Solution
10:  PerformTransition( $\alpha_t(\mathcal{SD})$ );
11: function PERFORMTRANSITION( $\alpha_t$ )
12:  for  $(sub^{et_j}, adv^{sd_k}) \in \alpha_t$  do
13:    if  $sub^{et_j}.sd_{previous} = \emptyset$  then
14:      Immediate_Transition( $sd_k$ );
15:    else
16:      Seamless_Transition( $sd_{previous}, sd_k$ )

```

To match advertisements to subscriptions, the function AssignDataSource looks for potential candidates for each event type in the list of related advertisements. In this regard, a data source announced by an advertisement is examined by a comparison between its current quality characteristics and the current quality policies' thresholds related to the subscription (e.g., the current accuracy level of a location data source should meet the policy of an accuracy level of fewer than 2 meters). This function assesses how many related quality policies can be fulfilled by this data source. If it meets all related quality policies (i.e., using the MeetAllQP function), the corresponding advertisement will be considered a qualified candidate for this subscription.

The costs of applying this sensing deployment include systems-side and monitoring costs which are calculated for each candidate and paired with its advertisement to form members of a list showing the eligible sensing deployments for each subscription (i.e., \mathcal{SD}). Then, a heuristic approach is applied on \mathcal{SD} considering satisfying the constraints in \mathcal{G} (i.e., HACS), to realize an approximate solution for the current situation. According to this newly generated solution, if the previous sensing deployment is changed, the transition between data sources can be done in two ways using the function PerformTransition.

The *immediate* transition model occurs when the previous sensing deployment is not available anymore. So, the controller should perform the transition as fast as possible with minimum delay, e.g., the sensor disconnected from the network. On the other hand, in *seamless* transition the previous data source is still available. Therefore, the controller performs the transition smoothly. For this type of transition, AQuA-CEP will process both data streams from previous and current producers concurrently in a period of β seconds in which the transition is happening from its invocation to its completion, e.g., when the sensor data

quality does not meet the quality policy's threshold. In some cases, the event stream's quality is degraded, and the stream is notably unreliable. To prevent wrong adaptation decisions, AQuA-CEP pauses monitoring and performs immediate transition.

4.3.3 Adaptation Strategies. Algorithm 3 shows the capabilities of the AQuA-CEP to adapt dynamically to the changes in the environment or process the queries based on the dynamic quality policy thresholds. Each QMA alarm has attributes such as type, corresponding sensing deployment (i.e., sd), quality policy (i.e., qp), and query identifier (i.e., Qid).

Algorithm 3 Alarm Processing

```

1: function PROCESSALARM( $A$ )
2:   switch ( $A.type$ ) do
3:     case SDUnavailability:
4:       for  $adv^{sd_i} \in ADV$  do
5:         if  $A.sd == sd_i$  then
6:            $ADV \leftarrow ADV - \{adv^{sd_i}\}$ ;
7:           AssignDataSource( $ADV, SUB$ );
8:     case ReducedQuality:
9:       Wait-Monitor( $A.sd$ );
10:    if  $A.sd$  Not Recovered then
11:      for  $adv^{sd_i} \in ADV$  do
12:        if  $A.sd == sd_i$  then
13:          Update( $adv^{sd_i}$ )
14:          AssignDataSource( $ADV, SUB$ );
15:    case ChangedQualityThreshold:
16:      for  $sub^{et_i} \in SUB$  do
17:        if  $A.qp \in QP_{sub^{et_i}}$  then
18:          Update( $sub^{et_i}$ )
19:          AssignDataSource( $ADV, SUB$ );
20:    case QueryEnded:
21:      for  $sub^{et_i} \in SUB_{A.Qid}$  do
22:         $SUB \leftarrow SUB - \{sub^{et_i}\}$ ;
23:        AssignDataSource( $ADV, SUB$ );

```

Upon the arrival of a QMA alarm, if the alarm's type indicates that the connection to a data source is lost and this sensing deployment is not available anymore (i.e., SDUnavailability), the controller removes all the related advertisements and performs data source re-assignment using the global optimizer. The next type of alarm is triggered by a reduction in the quality of data streams concerning the current thresholds of quality policies (i.e., ReducedQuality), e.g., when an obstacle blocks part of a motion detector's vision. In this case, the system performs a Wait-Monitor procedure in a specific period, in which AQuA-CEP checks the quality of produced events. If the data source can recover from this situation timely, our mechanism will continue with the current sensing deployment. Contrarily, if the lack of sufficient quality remains for the event stream, firstly the related advertisements to this sensing deployment will be updated with the new quality characteristics and then, a re-assignment

procedure will start. The main goal of the Wait-Monitor procedure is to prevent oscillation between data sources since it will lead to more switching costs and might produce a worse global solution.

Since the consumer is able to adjust the quality policy threshold during the runtime, various ranges are possible for thresholds according to the query model. Upon this type of variation, an alarm is triggered (i.e., ChangedQualityThreshold) to indicate that a new threshold should be taken into consideration. Hence, each subscription related to the changed quality policy has to be updated, and a new global re-assignment should be performed. Finally, if a query is finished on time or even ahead of time manually, the corresponding subscriptions will be removed from the set of subscriptions. In addition, the producers and CEP operators should disconnect from each other. Since the absence of those subscriptions may change the global solution, it is necessary to execute the AssignDataSource function again on the current available advertisements and subscriptions.

5 EVALUATION

In this section, we experiment with different ways of monitoring the quality of event detection and its corresponding adaptation strategies. The main goals of the evaluation are to figure out 1) does the data source assignment and switching satisfy the consumer's quality requirements while improving QoS, and 2) what are the cost or limitations involved in performing a transition among sensing deployments.

5.1 Simulation Setup

To run our scenarios, we created a Virtual Machine (VM) in *Oracle VM Virtual Box Manager* in which we installed *Ubuntu version 20* OS. We allocated 6 CPU cores with 100 percent execution capacity and 24 GB of main memory to the VM. We run complex event processing with multi-threading in this machine and create a thread for each of the issued queries, thereby, we could manage them simultaneously using Java.

For publish/subscribe communication of AQuA-CEP, we build on Apache Kafka as a distributed platform. Furthermore, for detecting complex events we build on *Flink-CEP* [7], which is a library implemented on top of Apache Flink. In our simulation, a Kafka server acts as an event broker that serves both data events and control events, as depicted in Figure 3. We monitor the quality of produced event streams in the QMA, which is located as an Apache Flink operator using the parallel QMA configuration model. In order to exhibit the potential of AQuA-CEP, we evaluated our mechanism in two different scenarios, one with the static quality policy and another with the dynamic quality policy.

5.2 Static Quality Policy Scenario

We analyze the performance of AQuA-CEP in terms of the event loss rate since this quality metric is directly influenced by assumed dynamics. We compared our approach with two baseline mechanisms. The first approach

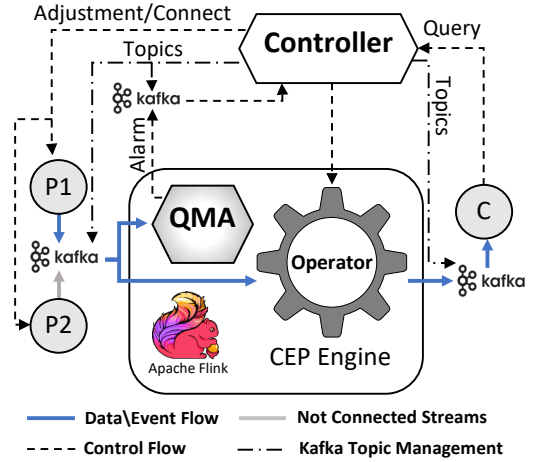


Figure 3. The evaluation model of AQuA-CEP.

is called *Optimal Dynamic Loss Rate (ODLR)*, which selects the best data source in terms of Loss Rate at the start of processing a query for each event type. Then, once the monitored loss ratio for produced event stream in runtime falls under its predefined characteristic, the controller first updates this feature with the new assessment and then performs a reassignment check. If this process results in opting for another data source, the controller will switch to the best producer in terms of the event loss rate. The second approach is *Optimal Static Loss Rate (OSLR)* which selects the best data source in terms of the event loss rate and stays with it until the end of the query.

To make our motivation clear in this scenario, we began with a *coverage problem* scenario [25], in which the objective is to solve the sensing placement problem to maximize the coverage of m important points using n sensors. To properly capture the capabilities of AQuA-CEP, we modified the coverage problem in a way that sensors are already located in the environment and their location cannot be changed. Moreover, instead of a set of m critical points, we have a dynamic group of queries to be answered. The goal is to optimize the total energy consumption by activating the set of sensors that can cover all queries.

```

1 SELECT event.*
2 FROM
3     // Selecting event stream
4     SELECT ds.stream
5     FROM DS
6     PATTERN
7     lossrate < QualityPolicy.threshold
8     Within window_size
9     WHERE
10     ds.type = 'Temperature' AND
11     ds.coverage(target_location) = True
12 WHERE event.value > Query.value
    
```

Listing 1. Applied query with static quality policy.

An example of a continuous query is shown in Listing 1, which aims to collect the temperature in a specific location accompanied by a static quality policy concerning a threshold value for the event loss rate characteristics of the sensing deployment. We applied variations in the event loss rate of all sensing deployment that represents the real

situations that might happen in the environment. We applied the same query with different values for thresholds on each approach with the same simulation setup.

The temperature data in [3] is used which contains two datasets and we chose one of them which has sensory data about temperature, pH, and turbidity from 30 cm below the water surface. We utilized this dataset because it has a real-world distribution of temperature data that makes our calculation more realistic. In our mechanism, each functioning sensor takes the data from this dataset and transmits its own transformed data according to its pre-defined quality characteristics. It means that each sensor will produce a unique temperature data stream according to its own features.

In addition, we estimated the amount of energy consumed by each sensor as described in [11]. This includes the energy for sensing the phenomena, processing the measurement, logging (i.e., reading data and writing it into the memory), communication, and transient energy (i.e., the transition energy to go from the idle state to the active state, and vice versa). We presumed that the distance between sensors to the gateway is the same, and they transmit packets of the same size.

5.3 Dynamic Quality Policy Scenario

We set side by side AQuA-CEP with *Optimal Dynamic Accuracy (ODA)* and *Optimal Dynamic Energy (ODE)* approaches. The former is pretty identical to the *ODLR* approach in a static quality policy scenario which always picks the best sensor in terms of accuracy and adjusts the sensing deployment regarding the person's current location. The latter assesses the energy consumption of the sensors as the distinguishing factor to opt for the best option on every occasion.

In this section, we introduce an IoT scenario that demonstrates the applicability of AQuA-CEP in the context of IoT surveillance. Most of the DCEP systems fail to support and utilize multiple sensing deployments in order to maintain a certain level of quality and react properly to the scenario dynamics, e.g., exceeding the coverage of a sensor due to the user's mobility. It indicates that DCEP systems require to rethink of switching between data sources at runtime. In order to demonstrate the use of quality policies, we consider a continuous query that aims to detect and warn any person approaching a Dangerous Area (DA) in an industrial zone depicted in Figure 4. In other words, if the person is located less than a threshold close to the dangerous area (i.e., within the Alarm Region (AR)), an alarm is triggered to inform him on his device. In our use case, the specified dynamic quality policy is "the quality of the person's location can degrade, as the person moves away from the borders of the dangerous area, but it should be as much accurate as possible when it is in a close proximity of the target location". In this policy, the quality metric is the location *accuracy* and the second parameter is the target's *location* (i.e., the accuracy level is determined according to the target's current location).

In this scenario, multiple entities are involved in the system including a CEP engine, stationary sensors, mobile sensors, people (i.e., targets), a consumer (i.e., query




Figure 4. Security monitoring: An example of sensing deployment transition using dynamic quality policy.

issuer), and guards that act cooperatively to protect the dangerous area. At first, the consumer in the factory needs to submit a query to control people's access to the dangerous area. To do so, the query should consist of the dangerous area's details and the alarm region. Besides, as in this case AQuA-CEP employs attribute-based access control, the query must specify the access policy tree of authorized targets. Once the required information is determined, the query will be submitted to AQuA-CEP.

Since a publish/subscribe mechanism is used for the management of people in this factory, a notification event of prohibition to approach a dangerous area is generated and all targets within the factory will be notified of the prohibition. That's because all targets in AQuA-CEP have already subscribed to these types of notification events when performing the admission process. Moreover, during the admission process, each target explains its attributes including its identity, and role in the factory. Also, it describes its carrying devices with their sensing capability (e.g., smartphones, tablets, smartwatches, wearables, etc.). In addition, it should be clearly detailed what type of data each device can produce and what sensing and communication technologies it can provide. Also, we assume that each target will give continuous access to their registered sensing deployments and not deliberately block the connection.

AQuA-CEP can access a sensing infrastructure based on a multitude of sensors deployed on a target's devices (e.g., smartphones), such as Bluetooth Low Energy (BLE), LTE, WiFi, and RFID sensors. Moreover, the system can also benefit from other types of positioning infrastructure embedded in the environment like cameras. Each of these sensing deployments has its own characteristics and in order to estimate the Energy Consumption (EC) in this scenario, we reuse the measurements collected from [8, 15, 23, 28, 32], as indicated in Table 1. Among all sensors, only *Camera* does not utilize the mobile phone's battery since it is a separate camera placed on the wall. Therefore, we assumed its energy consumption was equal to placing the phone in airplane mode.

The query correlates specific conditions - the target's position, the boundaries of the dangerous area, the alarm region, and the authorization status of the target - to detect a complex event of dangerous area violation. For this case, the perception of a location event could use different sensors to get an approximate position with the required

Table 1. Sensing Configurations and Their Characteristics

Name	Range (m)	EC (mW)	Accuracy (m)
BLE	70 - 100	426	1 - 3
RFID	1 - 12	375	0.1 - 2
WIFI	50 - 100	817	1 - 5
Camera	N/A	374	< 1
LTE	> Km	1634	< 1

quality level specified in the query's quality policies while optimizing the energy consumption of the application installed in the smartphone. Dependent on various aspects like coverage, location uncertainty, and sensing frequency, various query models can be utilized that meet the requirements of the trade-off between quality and energy efficiency. In this specific use case, since we do not need a very accurate position when the person is far from the dangerous area, AQuA-CEP uses the sensor which first meets the quality level of the query and then has the smallest amount of energy consumption for the smartphone. Here, the event quality becomes less precise, the farther a person is from the boundary of the dangerous area. Therefore, AQuA-CEP may change between different sensor deployments and appropriate query models.

Alike to the static quality policy scenario, we applied the query to the simulation environment multiple times considering the dynamic quality policy and analyzed the results to gain more insights. With a fixed route for a target person, we randomly generated coordination for DA and AR to evaluate the performance of all approaches when the DA is located close to different data sources.

```

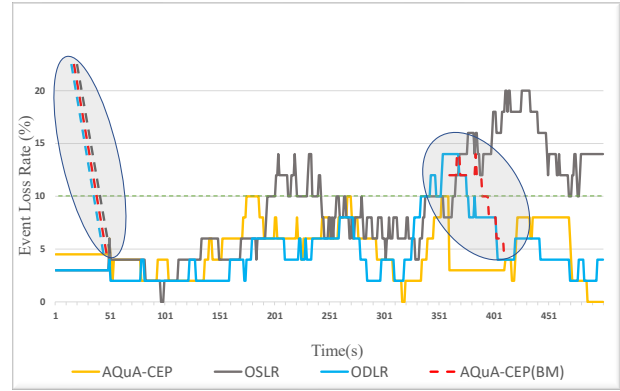
1 SELECT event.*
2 FROM
3     // Selecting Event Stream
4     SELECT ds.stream
5     FROM DS
6     PATTERN
7     Accuracy < CurrentQualityPolicy.threshold
8     Within window_size
9     WHERE ds.type = 'Location'
10    AND
11    // Selecting Current Quality Policy
12    CurrentQualityPolicy = (
13        SELECT qp
14        FROM Query.QP
15        WHERE qp.InRange(target_loc))
16 WHERE Distance(event.loc, Query.DA) < Query.AR

```

Listing 2. Applied query with dynamic quality policy.

The CEP-enriched SQL format of the applied query with a dynamic quality policy is represented in Listing 2, in which a quality policy is selected or changed when the target's current location is changed (i.e., using the `InRange` function). Once the current quality policy is specified, the candidate streams are chosen based on the predefined pattern.

Finally, a constraint satisfaction global optimization algorithm is hired to obtain an assignment solution for all queries from the determined candidate streams. In our

**Figure 5.** Event loss rate ratio during the execution for query 1 with the static quality policy.

simulation, we chose to apply *Choco-solver* [21], an open-source Java library for constraint programming that suits to satisfy our requirements.

5.4 Simulation Results

In this section, we will analyze our findings compared to the other two mentioned baseline strategies in both categories of static and dynamic quality policies. For the former, we opt for the event loss rate as a quality criterion, and for the latter, accuracy is the requested quality metric.

5.4.1 Static Quality Policy. The results regarding the static quality policy scenario have been illustrated in Figure 5, and Figure 6. We executed the simulation ten times for each query and the results were approximately similar. To challenge our approach, we selected the results with AQuA-CEP's worst performance.

Regarding the event stream loss rate in Figure 5, the chart displays the event loss rate for each approach in one execution. It can be seen that the OSLR approach shows the worst performance and proves the idea that each mechanism requires adapting to the dynamics. Both AQuA-CEP and ODLR select those sensing deployments meeting the threshold, which is illustrated as a green dashed line. Only two times, our approach exceeds the threshold of the event loss rate highlighted by the gray circles.

In the first quality policy violation, at the start of query execution, since we are employing a 50-sec window size over the event stream, we must wait for the window to be completed. So, we can not rely on the monitoring results, and we called this period as *Blind Monitoring (BM)* period. Such a period started once we chose a new sensing deployment. We masked this period with data from the characteristic of the new sensor for all three approaches and showed the simulation values with dashed lines. The second gray oval indicates the sensing deployment switching time for AQuA-CEP. Again, we showed the simulation results with the red dashed line. We cannot rely on this window information because the window comprises events that partially belong to the previous sensors, and the rest belongs to the newly selected sensor until the BM period is ended. Within this period, we can not perform




Figure 6. Total energy consumption for different sets of queries with the static quality policy.

adaptations since the outputs are not fully reliable. Therefore, the event loss rate results can go above the threshold and no sensing deployment switching would be triggered.

Hence, a lower number of BM periods results in a higher percentage of query duration being monitored. Having this fact in mind, AQuA-CEP achieves better results than the ODLR since it has fewer switching counts. The difference between AQuA-CEP and the ODLR has escalated considerably once more queries are submitted. That means the number of BM periods is increased significantly, leading to less reliability in quality monitoring that also influences the adaptation decisions. In addition, a higher number of sensing deployment switching results in taking more actions to activate or deactivate the sensors, stop analysis and monitoring on the previous sensors and start the procedure over the new sensing configuration.

From the energy consumption point of view, there is also a remarkable dissimilarity between these two approaches exhibited in Figure 6. The graph shows that the amount of energy consumed in AQuA-CEP is less than the ODLR approach. From the set of four queries onward, the total consumed energy seems equal. But, the reason is one or more queries in ODLR stopped processing, thereby, the consumed energy for them was equal to zero. On the other hand, queries in AQuA-CEP keep being answered until the end of simulation time. For a clearer comparison, we illustrated the results for both approaches until 250 sec of execution, when all queries are still active. Therefore, the results now give more insight into the energy consumption performance and are more likely to be comparable. Hence, the consumed energy for all sensors in AQuA-CEP is dramatically lower than the ODLR approach, respectively, proving the ability of our approach to optimize the server-side costs. It also can be seen that the difference between the two approaches is increasing by involving more queries in the execution which ascertains the fact that AQuA-CEP is more reliable than ODLR to deal with involving more queries.

5.4.2 Dynamic Quality Policy. Figures 7 and Figures 8 exhibit the simulation's outcome in the case of applying dynamic quality policy over the query processing. With energy consumption in mind as the comparison parameter, one can observe from Figure 7 that there is a dramatic




Figure 7. Total energy consumption for different sets of queries enriched by the dynamic quality policy.




Figure 8. Number of detected events and the summation of FP and FN for different sets of queries enriched by the dynamic quality policy.

distinction between AQuA-CEP and ODA. The rationale is ODA prefers the LTE sensor at all times for location tracking since it delivers the most accurate results while consuming the highest amount of energy among all sensing deployments. From this point of view, it can be concluded that employing the ODA mechanism can quickly lead to the phone's battery exhaustion while AQuA-CEP is able to preserve the battery at a satisfactory level. On the other hand, although the ODE approach is assumed the optimal approach in terms of energy, it consumes slightly less energy than AQuA-CEP. That's why we can claim the performance of AQuA-CEP in regard to energy consumption is near-optimal. It should be noted that the discrepancy between the results of these two techniques is raised slightly by initiating more queries, but it is still negligible.

In event-based systems, one of the main criteria to compare approaches is the number of False Positives (FP) and False Negatives (FN) in detecting complex events. In our example, an FN denotes a violation by entering a red-flagged area that occurred in the real world, but the event processing system could not catch it and triggered an alarm. Besides, an FP indicates a wrong violation that has been detected by the system while it is against the ground truth. Since both of these errors are feasible in our use case with small counts, we form a single number of their

summation that makes the differences more distinguishable as illustrated in Figure 8. Since ODA permanently answers queries with the most accurate sensors, it serves better than other methods. It could catch all the complex events without any FP or FN, but with the cost of draining the phone's battery. With fewer queries, ODE acts nearly the same as AQuA-CEP. But once more queries are involved, the situation becomes worse for the ODE approach by creating more FPs and FNs. In other words, the enlargement in the number of FNs and FPs is more evident in the bars related to ODE, while this summation in our results remains the same after adding more queries. It probably happens because of the lower detection capacity of data sources with the smallest energy consumption when the dangerous area is located in their closed vicinity. The sensors with less energy consumption level have less room for consumers and upon concurrent usage might deliver less accurate data leading to more FPs and FNs.

Similar to FN and FP in stream processing, *F-score* is a well-known performance measure in machine learning approaches that combines the other two measures, *precision* and *recall*, which are mostly employed to distinguish between classifiers [16] in terms of accuracy. Therefore, F-score can be used in stream processing to compare mechanisms in terms of accuracy. Analyzing the F-score shows an *ascending trend* in the reports which are approximately 0.947, 0.941, 0.961, 0.975, and 0.977 for sets of 1, 2, 3, 4, and 5 queries, respectively. While, the outcome for ODE as 1, 0.969, 0.961, 0.962, and 0.953 displays a *descending trend*. This proves the ability of AQuA-CEP to deal with involving more queries while maintaining the accuracy of event detection.

6 RELATED WORK

Reacting to environmental dynamics is highly important while using DCEP systems in IoT scenarios for maintaining the QoR and QoS at a satisfactory level based on consumers' quality requirements. To select the suitable sensing deployments for each query, AQuA-CEP performs a global optimization in sensing deployment configuration that takes into account the consumer-side constraints as their quality requirements and monitors the quality of produced events to assess and make adaptation decisions to maintain the quality of the produced outcomes. In this section, we compare our work to the related work in two key areas, sensor selection, and quality monitoring.

6.1 Sensor Selection

In the context of IoT networks, the Sensor Selection Problem (SSP) is a leading research direction to select the best set of sensors to achieve energy efficiency due to power limitations in sensor nodes [10, 30, 31]. More precisely, the selection in SSP-oriented research works is performed by picking a set of homogeneous sensors and executing the sensor aggregation. In the context of stream processing (e.g., DCEP systems), most runtime adaptive approaches concentrate on the operator networks, including adaptation on topology, deployment, processing, overload, fault tolerance, and infrastructure [6]. There are a few

approaches similar to AQuA-CEP which are called *data source switching* mechanisms [5, 14] and most of them only applied to video streaming applications [22].

Although both of the mentioned related approaches study the dynamic selection of sensors in an IoT environment to optimize the QoR, such optimizations are performed respecting specific data attributes or a set of specific fused sensor sources. However, integrating such methods in the context of DCEP requires linking them dynamically to different configurations of heterogeneous sensors. Only in this way, the flexibility of current DCEP systems in reconfiguring and rewriting the detection logic of complex events can be used to optimize for QoR.

6.2 Quality Monitoring

In DCEP systems, quality assessment has been mostly studied in the placement of detection logic over the available computing resources (e.g., [19]), and only a few research works focus on the adaptation of sensing deployment to react dynamics. Although these mechanisms are quite similar in performance to our proposed approach, they are focused only on one aspect of the system (e.g., CEP query language in [29]).

In the IoT environment, the *service composition* mechanisms consider the sensory data streams as services provided by the connected objects to be analyzed and deliver results to the corresponding applications and allow the interaction between consumers and smart objects of IoT environment [2, 9, 13]. Considering the vulnerability of IoT service quality to environmental dynamics, service composition techniques try to specify a set of quality metrics to analyze the quality of delivered streams to target applications. By employing heuristic and meta-heuristic techniques (e.g., [24]), several approaches attempt to find a global service composition solution while fulfilling QoS demands. These mechanisms tend to monitor and assess the quality of IoT services and adapt the system to maintain the quality, e.g., by training the Hidden Markov Models (HMM) to predict QoS. However, their quality expressivity is limited to defining the static thresholds causing inflexibility in acquiring more complex consumers' quality requirements. Moreover, AQuA-CEP is more efficient in terms of energy consumption since it endeavors to minimize the number of active sensing deployments and eventually fulfills the requested quality requirements.

7 CONCLUSION AND FUTURE WORK

In this work, we proposed AQuA-CEP which represents how to enable dynamic adaptation of sensing deployment configuration while observing the quality of produced events *and* their data sources. In addition, by proposing optimization criteria for the dynamic activation of sensors, our mechanism can help save resources in the sensing infrastructure. Our evaluation results demonstrated that AQuA-CEP outperforms two baseline approaches in terms of switching counts between sensing deployment and performed near-optimal in regard to the total energy consumed by the sensing network in static quality policy scenario. Moreover, by applying a dynamic quality policy,

AQuA-CEP achieved near-optimality in terms of energy consumption and quality measured in form of FP and FN. Moreover, the F-score results proved that AQuA-CEP has sufficient capabilities to fulfill consumers' quality requirements when more queries are involved.

In our future work, we will consider *priority* in the quality policy definition and investigate its impacts to support concurrent query processing. Estimating the switching overhead is another point of interest that requires further research. In addition, minimizing the blind monitoring periods can be attainable by predicting the data source switching time. We believe, building on a statistical analysis of the data sources' performance will be a promising direction. Finally, we plan to extend our proposed research by taking into account dynamic factors like quality degradation of data sources over time in dynamic quality policies.

References

- [1] K. Aberer, M. Hauswirth, and A. Salehi. 2006. *The Global Sensor Networks middleware for efficient and flexible deployment and interconnection of sensor networks*. Technical Report.
- [2] A.N. Abosaif and H.S. Hamza. 2020. Quality of service-aware service selection algorithms for the internet of things environment: A review paper. *Array* 8 (2020), p. 100041.
- [3] A.I. Arafat, T. Akter, M.F. Ahammed, M.Y. Ali, and A.A. Nahid. 2020. A dataset for the Internet of Things based fish farm monitoring and notification system. *Data in Brief* 33 (2020), p. 106457.
- [4] P. Arcaini, E. Riccobene, and P. Scandurra. 2015. Modeling and Analyzing MAPE-K Feedback Loops for Self-Adaptation. In *Proceedings of 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*. IEEE/ACM, pp. 13–23.
- [5] P. Asghari, A.M. Rahmani, and H.H.S. Javadi. 2018. Service composition approaches in IoT: A systematic review. *Journal of Network and Computer Applications* 120 (2018), pp. 61–77.
- [6] V. Cardellini, F. Lo Presti, M. Nardelli, and G.R. Russo. 2022. Runtime Adaptation of Data Stream Processing Systems: The State of the Art. *ACM Computing Surveys (CSUR)* 54, 11s (2022), pp. 1–36.
- [7] Apache Flink Community. 2017. *FlinkCEP*. <https://nightlies.apache.org/flink/flink-docs-master/docs/libs/cep/>.
- [8] L. Corral, A.B. Georgiev, A. Sillitti, and G. Succi. 2013. A method for characterizing energy consumption in Android smartphones. In *Proceedings of the 2nd International Workshop on Green and Sustainable Software (GREENS)*. IEEE, pp. 38–45.
- [9] F. Gao and E. Curry. 2020. Quality of Service-Aware Complex Event Service Composition in Real-time Linked Dataspaces. In *Real-time Linked Dataspaces*. Springer, pp. 169–190.
- [10] Y. Gao, M. Diao, and T. Fujii. 2019. Sensor Selection Based on Dempster-Shafer Evidence Theory under Collaborative Spectrum Sensing in Cognitive Radio Sensor Networks. In *Proceedings of 90th Vehicular Technology Conference (VTC2019-Fall)*. IEEE, pp. 1–7.
- [11] M.N. Halgamuge, M. Zukerman, K. Ramamohanarao, and H.L. Vu. 2009. An Estimation of Sensor Energy Consumption. *Progress In Electromagnetics Research B* 12 (2009), pp. 259–295.
- [12] I. Kolchinsky and A. Schuster. 2018. Efficient Adaptive Detection of Complex Event Patterns. *VLDB Endowment* 11, 11 (2018), pp. 1346–1359.
- [13] Ş. Kolozali, M. Bermudez-Edo, N. Farajidavar, P. Barnaghi, F. Gao, M.I. Ali, A. Mileo, M. Fischer, T. Iggena, and D. Kuemper. 2018. Observing the Pulse of a City: A Smart City Framework for Real-Time Discovery, Federation, and Aggregation of Data Streams. *Internet of Things Journal* 6, 2 (2018), pp. 2651–2668.
- [14] S.K. Lee, S. Yoo, J. Jung, H. Kim, and J. Ryoo. 2015. Link-Aware Reconfigurable Point-to-Point Video Streaming for Mobile Devices. *ACM Transactions on Multimedia Computing, Communications, and Applications* 12, 1 (2015), pp. 1–25.
- [15] C.T. Li, J.C. Cheng, and K. Chen. 2020. Top 10 technologies for indoor positioning on construction sites. *Automation in Construction* 118 (2020), p. 103309.
- [16] L. Li, B. Zhong, C. Hutmacher Jr, Y. Liang, W.J. Horrey, and X. Xu. 2020. Detection of driver manual distraction via image-based hand and ear recognition. *Accident Analysis & Prevention* 137 (2020), p. 105432.
- [17] M. Luthra, S. Hennig, K. Razavi, L. Wang, and B. Koldehofe. 2020. Operator as a Service: Stateful Serverless Complex Event Processing. In *Proceedings of International Conference on Big Data (Big Data)*. IEEE, pp. 1964–1973.
- [18] M. Luthra and B. Koldehofe. 2019. ProgCEP: A Programming Model for Complex Event Processing over Fog Infrastructure. In *Proceedings of the 2nd International Workshop on Distributed Fog Services Design*. ACM, pp. 7–12.
- [19] M. Luthra, B. Koldehofe, P. Weisenburger, G. Salvaneschi, and R. Arif. 2018. TCEP: Adapting to Dynamic User Environments by Enabling Transitions between Operator Placement Mechanisms. In *Proceedings of the 12th International Conference on Distributed and Event-based Systems (DEBS)*. ACM, pp. 136–147.
- [20] M. Nardelli, V. Cardellini, V. Grassi, and F.L. Presti. 2019. Efficient Operator Placement for Distributed Data Stream Processing Applications. *IEEE Transactions on Parallel and Distributed Systems* 30, 8 (2019), pp. 1753–1767.
- [21] C. Prud'homme and J.G. Fages. 2022. Choco-solver: A Java library for constraint programming. *Journal of Open Source Software* 7, 78 (2022), p. 4708.
- [22] C. Qin, H. Eichelberger, and K. Schmid. 2019. Enactment of adaptation in data stream processing with latency implications—a systematic literature review. *Information and Software Technology* 111 (2019), pp. 1–21.
- [23] R. Rieseboos, V. Degeler, and A. Tello. 2022. Smartphone-Based Real-Time Indoor Positioning Using BLE Beacons. In *Proceedings of 18th International Conference on Automation Science and Engineering (CASE)*. IEEE, pp. 1281–1288.
- [24] S. Sefati and N.J. Navimipour. 2021. A QoS-Aware Service Composition Mechanism in the Internet of Things Using a Hidden-Markov-Model-Based Optimization Algorithm. *IEEE Internet of Things Journal* 8, 20 (2021), pp. 15620–15627.
- [25] S. Singh, S. Kumar, A. Nayyar, F. Al-Turjman, and L. Mostarda. 2020. Proficient QoS-Based Target Coverage Problem in Wireless Sensor Networks. *IEEE Access* 8 (2020), pp. 74315–74325.
- [26] A. Slo, S. Bhowmik, and K. Rothermel. 2020. HSPICE: State-Aware Event Shedding in Complex Event Processing. In *Proceedings of the 14th International Conference on Distributed and Event-based Systems (DEBS)*. ACM, pp. 109–120.
- [27] A. Slo, S. Bhowmik, and K. Rothermel. 2022. State-Aware Load Shedding From Input Event Streams in Complex Event Processing. *IEEE Transactions on Big Data* 8, 5 (2022), p. 1340–1357.
- [28] Y. Wang, S. Han, Y. Tian, C. Xiu, and D. Yang. 2020. Is Centimeter Accuracy Achievable for LTE-CSI Fingerprint-Based Indoor Positioning? *IEEE Access* 8 (2020), pp. 75249–75255.
- [29] P. Weisenburger, M. Luthra, B. Koldehofe, and G. Salvaneschi. 2017. Quality-Aware Runtime Adaptation in Complex Event Processing. In *Proceedings of 12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*. IEEE/ACM, pp. 140–151.
- [30] I. Younas and A. Naeem. 2022. Optimization of sensor selection problem in IoT systems using opposition-based learning in many-objective evolutionary algorithms. *Computers & Electrical Engineering* 97 (2022), p. 107625.
- [31] J. Zhang, J. Du, and L.R. Dai. 2021. Sensor Selection for Relative Acoustic Transfer Function Steered Linearly-Constrained Beamformers. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 29 (2021), pp. 1220–1232.
- [32] L. Zou, A. Javed, and G.M. Muntean. 2017. Smart mobile device power consumption measurement for video streaming in wireless environments: WiFi vs. LTE. In *Proceedings of International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*. IEEE, pp. 1–6.