

Detecting contaminants in smart buildings by exploiting temporal and spatial correlation

G. Boracchi, M. Michaelides, M. Roveri

Abstract—Monitoring the indoor air quality is one of the most critical activities within a smart building environment. The introduction of contaminant sources inside the building envelope can compromise the air quality and possibly endanger the lives of the inhabitants. In this paper, a new contaminant-detection system is proposed for the prompt and effective detection (and isolation) of contaminant sources. Specifically, we address the challenging scenario where the contaminant of interest is also naturally present in the indoor building environment (e.g. CO₂). A key feature of the proposed system is that it does not require a model of the contaminant propagation, but relies instead in its ability to exploit the temporal and spatial relationships present in the datastreams acquired by the sensors deployed within the smart building. The effectiveness of the proposed system has been evaluated on a reference testbed.

I. INTRODUCTION

Smart homes/buildings¹ represent an important and challenging application scenario for new Internet-of-Things (IoT) technologies and solutions. These cyber-physical systems that are typically composed of networks of units equipped with sensors and actuators interacting with the environment in which they operate, require intelligent mechanisms to improve, guarantee or maintain the application performance in real-world conditions. In a smart building application, sensors are typically used for measuring the temperature, the humidity and the concentration of various contaminants of interest (e.g., carbon dioxide), while actuators are envisioned for automatically opening the doors and windows, controlling the exhaust fans and the Heating, Ventilation and Air Conditioning (HVAC) units.

Within a smart building environment, a critical activity for the safety and comfort of the occupants is the assessment and control of the Indoor Air Quality (IAQ). Note that the IAQ can be easily compromised from a contaminant source (e.g., chemical, biological or radioactive substance) introduced inside the building envelope as a result of an accident (e.g., CO released from a faulty furnace) or a terrorist attack (e.g., anthrax). Moreover, the IAQ could be affected by the abnormally high concentration of a contaminant (i.e. CO₂) in a particular zone of the building due to increased occupancy and/or inappropriate ventilation conditions. Note that

G. Boracchi and M. Roveri are with the Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Milan, Italy, {giacomo.boracchi,manuel.roveri}@polimi.it; M. Michaelides is with the Department of Electrical Engineering, Computer Engineering and Informatics, Cyprus University of Technology, 30 Archbishop Kyprianos Str., CY-3036 Lemesos, Cyprus, {michalis.michaelides}@cut.ac.cy

¹In the sequel, the terms “smart homes” and “smart buildings” will be used interchangeably.

detecting this situation is particularly challenging, since the contaminant is also naturally present in the atmosphere. Under all the aforementioned safety-critical conditions, it becomes of paramount importance to promptly detect the presence of the contaminant source and estimate its location, in order to take the necessary actions for ensuring the safety of the occupants. These could involve simple actions like opening doors/windows or controlling the HVAC system, up to the complete evacuation of the building when this is deemed necessary.

In this paper, we target the challenging problem of detecting sources of contaminants and/or pollutants that are also normally present in the atmosphere (e.g., CO₂), thus also inside the building environment. To this purpose, we propose the Contaminant-Detection System (CDS) portrayed in Figure 1, which is meant for smart homes that are equipped with appropriate sensors for measuring the contaminant concentration at different locations. The stream of measurements from all these sensors are first gathered by an ad-hoc sensor network and then analyzed by the CDS. The key idea is to exploit both the temporal correlation that characterizes each stream of contaminant measurements, and the spatial correlation that characterizes streams of measurements acquired in different zones. These include rooms that are directly connected by doors, windows or air handling units, which exhibit causal relationships between their measurements. Thus, under nominal conditions, we learn the relationships characterizing the streams of measurements in these particular zones by means of Input-Output predictive models. During the contaminant-detection activity, we analyse the discrepancy (residuals) between the measured and the predicted output in each zone of the smart home to detect the presence of the contaminant sources. In particular, contaminants are detected through a suitable Change-Detection Test (CDT), which assesses changes in the statistical behaviour of the residuals. When a contaminant event is detected and validated (through an additional statistical analysis), the zone containing the contaminant source is isolated and emergency procedures can be activated.

Examples of CDSs in the literature can be found in [1], [2]. In their simplest form, these systems are composed of a triggering mechanism, which detects contaminants as soon as the concentration measurement of any sensor exceeds a threshold. This threshold is usually defined to allow the CDS to operate at a pre-specified false alarm rate based on the sensor noise or the natural fluctuation of the contaminant, both of which are often difficult to model/estimate in real-world scenarios. Therefore, these systems are limited in sensitivity

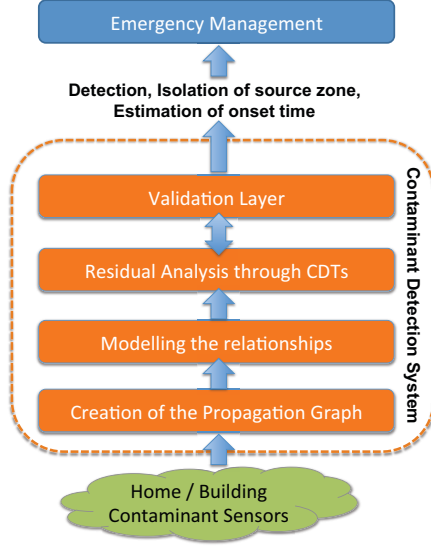


Fig. 1. Overview of the Contaminant Detection System.

by the choice of the threshold, which is often chosen large enough to avoid frequent false alarms. Note that frequent false alarms can lead to the occupants’ reluctance to follow the emergency procedures with catastrophic sequences (e.g., the cry-wolf effect). More sophisticated CDSs employ a contaminant propagation model in an attempt to utilize all the available information coming from the building to improve the detection performance and isolate the contaminant source (i.e., estimate its location). These include the adjoint probability method [3], the state space method [4] and the Bayesian updating method [5]. They differ in the assumptions they make about the source characteristics, the noise and the uncertainty and the way they build the model. In particular, the Bayesian method implicitly “learns” a model by constructing a large-scenario database before the event. Note that, in a real-world setting, these models are very difficult to compute accurately because of the various environmental uncertainties involved in estimating the propagation of contaminants in indoor building environments (i.e., the wind velocity and direction, the temperature and the various leakages). To address this problem, a new generation of CDSs following a cognitive approach has been recently proposed in the literature [6]. Cognitive CDS (also referred to as intelligent CDS) are able to promptly detect variations in the contaminant under monitoring by using appropriate statistical tools (change-detection tests, change-point methods), as well as distinguish between the actual presence of the contaminant and a fault affecting one of the sensors deployed within the building.

The CDS proposed in this paper falls in this last category as it uses similar statistical tools for detection. It is fundamentally different in concept, however, from what was presented in [6] for the following reasons: (i) in nominal conditions, the concentration of the monitored contaminant can naturally follow a time-varying evolution, (ii) most relevant relationships

between pairs of sensors are identified by constructing a propagation graph from topological and airflow information (iii) Input-Output Predictive Models are used to model the temporal and spatial relationships characterizing the stream of measurements, and (iv) the contaminant is detected by monitoring the discrepancy between the acquired and the predicted measurements using the estimated relationships. A wide experimental campaign encompassing data from a well-know testbed (i.e., the Holmes’s house), validate the effectiveness of the proposed approach.

The paper is organized as follows: Section II states the problem; the proposed CDS for smart buildings is detailed in Section III; experiments are presented in Section IV and, finally, conclusions are drawn in Section V.

II. THE PROBLEM STATEMENT

Let us consider an intelligent building composed of N zones. Each zone is equipped with a sensor measuring the concentration of a specific contaminant of interest. At time T^* (onset time), a contaminant source is introduced at zone i^* with $1 \leq i^* \leq N$, which we denote as the *source zone*. Our goal is to process all the sensor measurements to promptly detect the presence of the contaminant source and to effectively estimate the onset time T^* and the source zone i^* .

We model the measurement $m_i(t)$ of sensor in zone i at time t as

$$m_i(t) = c_i(t) + \Delta_i(t) + \eta(t), \quad (1)$$

where $c_i(t)$ is the natural concentration of the contaminant present in each zone, $\Delta_i(t) \geq 0$ is the amount of contaminant introduced in the i -th zone by the contaminant source and $\eta(t)$ accounts for independent and identically distributed (i.i.d.) measurement noise. Note that in our previous work [6], $c_i(t)$ was assumed to be zero or constant. Here, we do not make any assumption on $c_i(t)$ (thus, $c_i(t)$ can be time-varying, e.g., a time-series).

We define $\tau_i \geq T^*$ as the time instant at which the contaminant first appears in the i -th zone, i.e.,

$$\begin{cases} \Delta_i(t) = 0, & t < \tau_i \\ \Delta_i(t) > 0, & t \geq \tau_i. \end{cases} \quad (2)$$

It should be noted that $\Delta_i(t) > 0$, could be attributed to two different reasons: (i) the contaminant source was inserted in zone i (and in this case, $\tau_i = T^*$) or (ii) the contaminant source was inserted in a zone j located “upstream” w.r.t. i , such that the i -th zone receives flow from the j -th zone (and in this case $\tau_i > T^*$).

In general, the propagation of a contaminant inside the building envelope depends on the airflows among the various building zones, which are influenced by a number of different factors including (i) the building structure (e.g., the interconnections of the various zones through doors and openings), (ii) environmental conditions (e.g., temperature, wind direction and velocity), (iii) and the HVAC operational mode (or any other type of fan causing a forced flow). For the purposes of this paper, we assume that these airflows can be measured

(or suitably estimated), which is a reasonable assumption in the context of smart buildings. In particular, we define a *flow matrix* accounting for the magnitude of the airflows among the different zones of the building:

$$Z = \begin{bmatrix} z_{1,1} & \cdots & z_{1,N} \\ \vdots & \ddots & \vdots \\ z_{N,1} & \cdots & z_{N,N} \end{bmatrix}, \quad (3)$$

where $z_{i,j} \in \mathbb{R}^+ \cup \{0\}$ indicates the magnitude of the flow from the i -th zone to the j -th zone, which can be positive when a flow exists or 0 otherwise; for this reason $z_{i,i}(t) = 0, i = 1, \dots, N$. An example of a flow matrix is given in Figure 3 corresponding to the Holmes's house case-study shown in Figure 2. Note, that in this paper, we are assuming steady-state airflow conditions, even though in the more general case, Z can also be made time-varying. Also, it is important to stress that the proposed CDS does not need to accurately know the flows, since it can operate provided rough estimates of the flow magnitude and direction.

III. THE PROPOSED INTELLIGENT CONTAMINANT DETECTION SYSTEM

The core of the proposed intelligent CDS lies in the ability to model the relationships among the streams of contaminant measurements acquired by the N sensors of the smart building. To achieve this goal, the proposed system relies on a training set of measurements from each sensor, which is arranged in a matrix

$$TS = \begin{bmatrix} m_1(1) & \cdots & m_1(t_0) \\ \vdots & \ddots & \vdots \\ m_N(1) & \cdots & m_N(t_0) \end{bmatrix}, \quad (4)$$

that is acquired without contaminant sources inside the building (i.e., $T^* > t_0$). First, the training set is used to estimate predictive Single-Input/Single-Output models between the streams acquired in interconnected building zones (i.e., where flow exists from one zone to the other). Then, the presence of contaminant sources is monitored in each zone by analyzing the discrepancy between what the sensor in the zone is measuring and the estimate provided by the corresponding predictive model. Finally, once a substantial increase in the discrepancy has been detected, a further level of analysis is activated to confirm the change, isolate the zone in which the contaminant has been inserted and estimate the onset time.

The proposed CDS is summarized in Figure 1, while the corresponding algorithm is detailed in Algorithm 1. Next, we explain each of the algorithm steps in detail.

A. Creating the propagation graph

The first step of the proposed CDS consists in analyzing the flow matrix Z to create a *propagation graph* that contains those pairs of zones that are the most relevant for building predictive models of the sensor measurements. To this purpose, we select pairs of building zones that are connected by a sufficiently large airflow. In particular, we first build a directed graph $\mathcal{G}^0 = \{\mathcal{N}^0, \mathcal{E}^0\}$ whose node set \mathcal{N}^0 contains all the N

INPUT: Flow matrix Z , Training set TS , Threshold Γ_f ; Creation of the dependency graph $\mathcal{G} = \{\mathcal{N}, \mathcal{E}\}$ based on Γ_f and the analysis of the flow matrix Z ;

```

for (Each arc  $e_{(i,j)} \in \mathcal{E}$ ) do
    Estimate  $f_{i,j}^\theta$  on the training set  $TS_{rel}^{(i,j)}$ ;
    Compute the training residual sequence
     $\{r_{i,j}(t) = t_s + 1, \dots, t_0\}$ ;
    Configure the CDT on the training residual sequence;
end
while (1) do
    Each unit: acquire  $m_i(t)$ ;
    for (Each arc  $e_{(i,j)} \in \mathcal{E}$ ) do
        Compute  $r_{i,j}(t) = m_j(t) - \widehat{m}_{i,j}(t)$ ;
        Run the CDT on  $r_{i,j}(t)$ ;
        if Detection at time  $\widehat{T}$  then
            Run the refinement procedure to estimate  $\widehat{T}^*$ 
            the time instant the change started;
            Run the statistical hypothesis test to validate
            the difference in the statistical behaviour
            between the training residual sequence and
             $\{r_{i,j}(t), t = \widehat{T}^*, \dots, \widehat{T}\}$ ;
            if (Detection is validated) then
                Isolated zone:
                 $\widehat{i} = \operatorname{argmax}_{k \in \{i,j\}} \sum_{t=\widehat{T}^*, \widehat{T}} m_k(t)$ ;
            end
        end
    end
end

```

Algorithm 1: The proposed Contaminant-Detection System for smart buildings.

sensors of the smart home and \mathcal{E}^0 is the set of all the arcs connecting the nodes according to the building propagation paths. Then, we remove arcs from \mathcal{E}^0 according to the following rule: the arc $e_{(i,j)}$ connecting sensor i to sensor j is removed when the flow from zone i to zone j falls below a user-defined threshold Γ_f , i.e., $z_{i,j} < \Gamma_f$. Next, we also remove all the isolated nodes from \mathcal{N}^0 . The resulting final propagation graph, which we define as $\mathcal{G} = \{\mathcal{N}, \mathcal{E}\}$, is made available to the next phases of the intelligent CDS.

An example graphically portraying the construction of a propagation graph is presented in Figure 2 for the Holmes's house case-study, while the corresponding flow matrix Z is shown in Figure 3 (here $\Gamma_f = 5$). From Figure 2(c), it becomes evident that only large-flow relationships are retained in the final graph, while nodes $\{6, 7, 10, 11\}$ cannot be monitored using the proposed method.

B. Modeling the Relationships Among Zones

For each arc $e_{(i,j)} \in \mathcal{E}$, we estimate $f_{i,j}^\theta$, an Input-Output predictive model in the form

$$\widehat{m}_{i,j}(t) = f_{i,j}^\theta(m_j(t-1), m_j(t-2), \dots, m_j(t-k_j), m_i(t), m_i(t-1), m_i(t-2), \dots, m_i(t-k_i)), \quad (5)$$

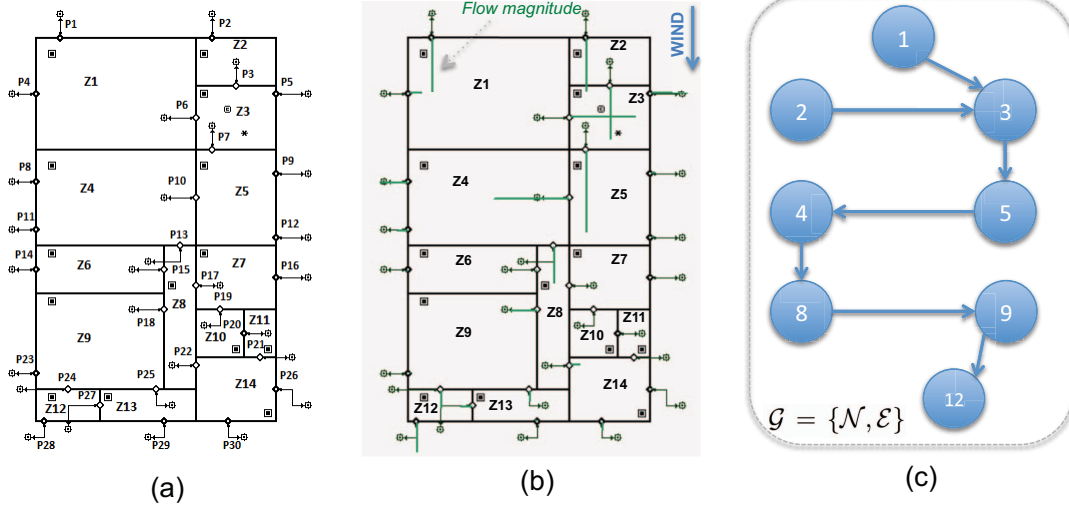


Fig. 2. The Holmes's house case study: from the building map to the construction of the propagation graph. (a) The map of the building; (b) The airflows between zones of the building when the wind direction is from North to South. The direction and the length of the green lines on the building map represent the direction and the magnitude of the respective airflows. The corresponding flow matrix Z is shown in Figure 3; (c) the associated propagation graph computed as described in Section III-A. Here, Γ_f has been set to 5.

	Z1	Z2	Z3	Z4	Z5	Z6	Z7	Z8	Z9	Z10	Z11	Z12	Z13	Z14
Z1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Z2	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Z3	22.0625	17.8777	0	0	0	0	0	0	0	0	0	0	0	0
Z4	0	0	0	0	25.6359	0	0	0	0	0	0	0	0	0
Z5	0	0	27.7099	0	0	0	0	0	0	0	0	0	0	0
Z6	0	0	0	0	0	0	0	0.3645	0	0	0	0	0	0
Z7	0	0	0	0	0	0	0	0.1836	0	0	0	0	0	0
Z8	0	0	0	12.7284	0	0	0	0	0	0	0	0	0	0
Z9	0	0	0	0	0	0	0	9.0531	0	0	0	0	0	0
Z10	0	0	0	0	0	0.0000	0	0	0	0	0	0	0	0
Z11	0	0	0	0	0	0	0	0	0	0.0000	0	0	0	0
Z12	0	0	0	0	0	0	0	0	5.9060	0	0	0	0	0
Z13	0	0	0	0	0	0	0	0	3.5566	0	0	0	3.9651	0
Z14	0	0	0	0	0	0	0	3.1273	0	0	0.0000	0	0	0

Fig. 3. The flow matrix Z for the Holmes's house case study shown in Figure 2(a) when the wind direction is from North to South.

where $\hat{m}_{i,j}(t)$ is the predicted value of the sensor j at time t , and f^θ is the predictive model having parameter set θ ; k_i and k_j represent the order of the input and output component, respectively. Note that f^θ can be linear (e.g., Auto Regressive with eXogenous input -ARX- models, Auto Regressive Moving Average with eXogenous input -ARMAX- models, Output Error -OE- models) or non linear (e.g., Non-linear ARX models, Hammerstein-Wiener models) [7], and that it takes as input also the measurements in zone i . Parameters θ are estimated from a training set $TS_{rel}^{(i,j)}$ that is composed as follows:

$$TS_{rel}^{(i,j)} = \begin{bmatrix} m_i(1) & \cdots & m_i(t_s) \\ m_j(1) & \cdots & m_j(t_s) \end{bmatrix}, \quad (6)$$

with $t_s < t_0$ representing the first t_s samples of the i -th and j -th row of TS .

The choice of the model family is application-dependent and might follow the intuitive principle where simple models (whenever effective) should be preferred to more complex ones (which generally require more samples to be effectively trained). The model orders can be either user defined or

automatically identified through a system identification phase: in the latter case $TS_{rel}^{(i,j)}$ will be further partitioned into two subsets for training and order estimation, respectively. As shown in the next section, the rest of the training set TS will be used to configure the CDT. The partitioning between training samples for model estimation and CDT configuration is meant to prevent overfitting, which might occur when the CDT is configured on data used for model estimation. For the specific scenarios investigated in this paper, we opted for linear input-output Auto-Regressive with eXogenous input (ARX) predictive models, as they revealed to be very effective in characterizing the relationships present in the acquired streams of measurements (a preliminary analysis of the fitness ability has been carried out).

For each estimated predictive model, we compute the *residuals*, as the discrepancy between the acquired contaminant measurement $m_j(t)$ and $\hat{m}_{i,j}(t)$, its prediction from $f_{i,j}^\theta$, i.e.,

$$r_{i,j}(t) = m_j(t) - \hat{m}_{i,j}(t). \quad (7)$$

These residuals measure the fitness of the estimated model

$f_{i,j}^\theta$ on each of the incoming data, and are monitored to detect contaminant sources.

C. Monitoring Residuals

When no contaminant sources are present in the building and the prediction provided by $f_{i,j}^\theta$ in (7) is accurate, the residuals $r_{i,j}(t)$ are expected to form an i.i.d. sequence. Therefore, we monitor the statistical behaviour of each residual sequence $r_{i,j}(t)$ by means of CDTs, namely sequential techniques to detect changes in datastreams. More specifically, we considered the ICI-based CDT [8] for its effectiveness and low computational load. The ICI-based CDT extracts Gaussian-distributed features from datastreams (i.e., the sample mean and a power-law transformation of the sample variance) and relies on the Intersection-of-Confidence-Intervals (ICI) rule to detect changes [9]. In this CDT, the trade-off between detection promptness and false-positive rate is regulated by the parameter Γ (large values of Γ reduce the false-positive rate at the expense of an increased detection delay, while small values of Γ reduce the detection delay at the expense of an increase in the false-positive rate). The ICI-based CDT associated to the prediction model $f_{i,j}^\theta$ is trained on the residual sequence

$$\{r_{i,j}(t), t = t_s + 1, \dots, t_0\}. \quad (8)$$

After the training phase (i.e., for $t > t_0$), the ICI-based CDT continuously inspects the stream of residuals $r_{i,j}(t)$ to detect changes in their distribution. Remember that a change in the residual indicates that the estimated model $f_{i,j}^\theta$ is no more able to accurately predict the measurements coming from zone i , and therefore this signifies a possible variation in the relationship due to the presence of a contaminant source inside the building. Note that the residuals (8) used for configuring the CDT are assumed to be computed in nominal conditions, hence they are representative of the residual distribution in stationary conditions.

D. Validation Layer

Assume that at time \hat{T} a CDT raises a detection. Then, change-validation procedures are activated to determine whether the detection corresponds to an actual change in the residual distribution or a false alarm. When the change is validated, the source zone and the onset time are estimated. To achieve this goal, the refinement procedure [8] of the ICI-based CDT is at first executed to compute \hat{T}^* , an estimate of the time instant when the change started in the sequence. The detection is then validated by an hypothesis test that determines whether the residuals between \hat{T}^* and \hat{T} have been generated from the same distribution that generated the residuals inside the initial training sequence. When the change is confirmed, all the measurements acquired after \hat{T}^* from the zones involved in the relationship detecting the change are further analyzed to isolate the zone in which the contaminant source has been inserted.

In more detail, let $e_{(\bar{i},\bar{j})}$ be the relationship yielding the stream of residuals where the change was detected. We then

compute an estimate \hat{T}^* of the time instant when the contaminant source has been inserted in the building (in either zone \bar{i} or \bar{j}), by the refinement procedure of the ICI-based CDT (Algorithm 3 in [8]), which is applied to the juxtaposition of a suitable buffer containing the most recent residuals of the stream $r_{(\bar{i},\bar{j})}$, and the initial ones. Alternatively, \hat{T}^* can be computed by an ensemble of change-point methods [10] to cope with correlation within the residual stream. Here, we are implicitly assuming that the first relationship detecting a change involves the source zone.

Afterwards, following the solution presented in [6], we consider a statistical hypothesis test (i.e., the Hotellings T^2 statistic [11] on the features extracted by the ICI-based CDT) to assess whether the features extracted from the initial training sequence (8) and from $\{r_{\bar{i},\bar{j}}(t), t = \hat{T}^*, \dots, \hat{T}\}$ have a different expectation, provided a predefined confidence level α (which is typically set to 0.01 or 0.05).

When the change in the residual is validated, the isolation phase is activated; otherwise, the change raised by the CDT is considered to be a false positive detection, hence discarded, and the contaminant-detection activity proceeds. In the isolation phase, the contaminant measurements in zones \bar{i} and \bar{j} after \hat{T}^* are inspected to estimate the zone \hat{i} in which the contaminant source has been inserted as follows:

$$\hat{i} = \operatorname{argmax}_{k=\{\bar{i},\bar{j}\}} \sum_{t=\hat{T}^*}^{\hat{T}} m_k(t). \quad (9)$$

In other words, the estimated isolation zone \hat{i} is the one (between \bar{i} and \bar{j}) characterized by the largest amount of accumulated contaminant after the estimated onset time \hat{T}^* .

IV. EXPERIMENTAL RESULTS

The effectiveness of the proposed CDS has been assessed on a dataset generated by the Matlab-CONTAM toolbox [12], that is a software specifically designed to simulate the presence of a contaminant in intelligent buildings. In particular, in our experiments, we considered the Holmes's house [13] depicted in Figure 2(a), which comprises of 14 zones (Z1 to Z14) as well as 30 leakage path openings corresponding to windows and doors (P1 to P30). The contaminant of interest (i.e., CO2) is present in the atmosphere with a mean concentration of $50g/m^3/h$ modeled as a pseudorandom sequence. The time interval between the transitions (jumps) follows a Markov process with 0.1 transition probability, while the magnitude of the sequence after a transition is a random number between $[45, 55]^2$.

In our simulations we consider two different scenarios characterized by wind speed of $10m/s$ and fully open leakage path openings; in particular

- *Scenario 1*: a contaminant source of emission rate 100 g/h is placed in Zone 5 with wind direction 0° at sample 2000;

²Note that the specific numbers are simply chosen for simulation purposes. In reality, we expect the contaminant of interest to be described by a stochastic process whose parameters will depend on the specific environment.

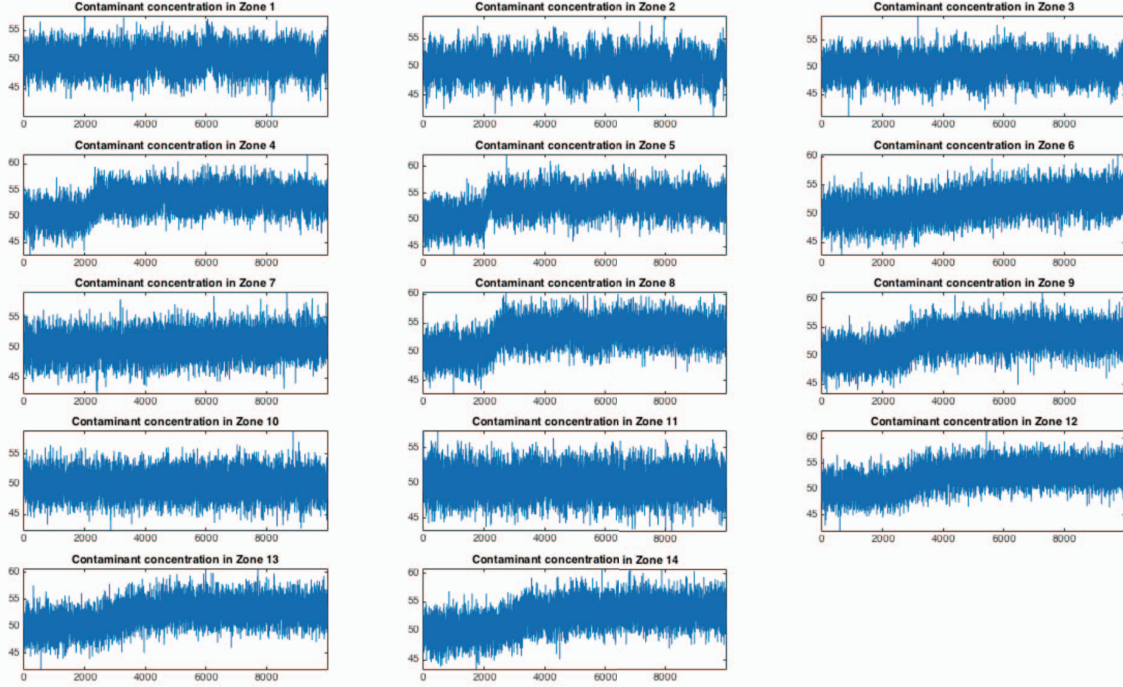


Fig. 4. An example of acquired measurements in all the 14 zones of the Holmes’s house in *Scenario 1* and with $\sigma = 2$.

- *Scenario 2*: a contaminant source of emission rate 100 g/h is placed in Zone 9 with wind direction 90° at sample 2000;

Each simulation lasts 10000 samples.

Following the model in (1), we assume that contaminant measurements are affected by random noise $\eta(t)$ distributed as a zero-mean *Gaussian* distribution $\mathcal{N}(0, \sigma^2)$, with σ ranging from 0.5 to 2. Furthermore, we assume there is a sensor in each of the 14 zones of the building. The considered predictive model $f_{i,j}^\theta$ is the linear ARX model, while the orders k_i and k_j have been experimentally fixed to 3 and 3, respectively. The parameter Γ of the ICI-based CDT ranges from 2 to 3.5, while t_0 and t_s are 1000 and 500 samples, respectively.

We emphasize that, in this experimental section, we significantly extend what was presented in [6] by considering a scenario where we have the natural presence of a contaminant within the building before the introduction of the contaminant source inside the building envelope.

We compare the proposed intelligent CDS against (i) a solution based on residuals computed from *AutoRegressive models (AR)*, i.e., models that rely only on data coming from the specific monitored zone and that do not leverage any relationship among different zones, and (ii) a solution based on *Exogenous-Input models (X)*, i.e., models similar to ARX but without the autoregressive component, thus relying only on data coming from a different zone to predict the contaminant in another zone. The isolation phase of the AR solution simply considers as source zone the first one detecting a change (there

is no ambiguity since it does not involve measurements from other zones). Differently, the X solution relies on the same isolation mechanism of the proposed CDS (Section III-D).

To compare the aforementioned solutions, the following figures of merit are computed over 250 runs:

- False Positive Rate (*FPR*), the percentage of experiments in which the change was detected before T^* ;
- False Negative Rate (*FNR*), the percentage of experiments in which the change was not detected;
- Detection Delay (*DD*), that is the average value (in samples) of $\hat{T} - T^*$ for changes correctly detected (i.e., when $\hat{T} > T^*$);
- ϵ_{iso} , the percentage of experiments in which the source zone has not been correctly estimated.

Figure 4 shows examples of the measurements in all the zones of *Scenario 1* when $\sigma = 2$, while Figure 5 depicts the computed residual for the relationship (*Input Z3-Output Z5*).

The experimental results that are presented in Table I and II, show the advantages provided by the proposed solution. By looking at Table I and II, several comments arise. First, the proposed system is truly able to promptly detect the occurrence of a contaminant source within the building for both considered Scenarios. In fact, the *FNR* is zero for all the configurations meaning that the contaminant source is always detected. Remarkably, this ability does not come at the expense of the *FPRs* that are close to zero in both the scenarios where $\Gamma > 2$. We also emphasize that the *DD* is relatively small meaning that the contaminant sources are promptly detected

Noise σ	Gamma	Proposed				Autoregressive (AR)				Exogenous Inputs (X)			
		FNR	FPR	DD	ϵ_{iso}	FNR	FPR	DD	ϵ_{iso}	FNR	FPR	DD	ϵ_{iso}
0.5	2.00	0.000	0.796	73.7	0.000	0.880	0.000	5604.0	0.833	0.000	0.944	68.6	0.000
	2.50	0.000	0.252	86.0	0.000	0.952	0.000	5200.0	0.750	0.000	0.400	77.2	0.007
	3.00	0.000	0.024	107.8	0.000	0.992	0.000	NaN	NaN	0.000	0.064	96.4	0.000
	3.50	0.000	0.000	138.7	0.000	0.996	0.000	NaN	NaN	0.000	0.016	124.9	0.000
1	2.00	0.000	0.120	116.4	0.000	0.876	0.008	2170.0	0.931	0.000	0.324	100.9	0.047
	2.50	0.000	0.032	144.9	0.004	0.968	0.000	3830.0	0.750	0.000	0.068	128.3	0.017
	3.00	0.000	0.004	184.9	0.000	1.000	0.000	NaN	NaN	0.000	0.008	165.4	0.000
	3.50	0.000	0.000	226.1	0.000	0.996	0.000	NaN	NaN	0.000	0.000	206.1	0.004
2	2.00	0.000	0.004	208.5	0.004	0.824	0.000	4864.0	0.886	0.000	0.060	187.1	0.030
	2.50	0.000	0.000	259.6	0.000	0.968	0.000	4666.7	0.625	0.000	0.012	236.3	0.012
	3.00	0.000	0.000	314.5	0.000	0.988	0.000	NaN	NaN	0.000	0.000	290.0	0.004
	3.50	0.000	0.000	366.1	0.000	1.000	0.000	NaN	NaN	0.000	0.000	342.2	0.000

TABLE I
EXPERIMENTAL RESULTS ON SCENARIO 1.

Noise σ	Gamma	Proposed				Autoregressive (AR)				Exogenous Inputs (X)			
		FNR	FPR	DD	ϵ_{iso}	FNR	FPR	DD	ϵ_{iso}	FNR	FPR	DD	ϵ_{iso}
0.5	2.00	0.000	0.628	116.5	0.117	0.828	0.004	5580.0	0.929	0.000	0.844	114.2	0.050
	2.50	0.000	0.136	133.3	0.009	0.944	0.000	6780.0	0.929	0.000	0.256	128.8	0.011
	3.00	0.000	0.012	150.3	0.000	0.996	0.000	NaN	NaN	0.000	0.036	143.1	0.008
	3.50	0.000	0.004	168.8	0.000	0.996	0.000	NaN	NaN	0.000	0.012	160.2	0.000
1	2.00	0.000	0.020	163.8	0.033	0.856	0.000	3900.0	0.861	0.000	0.128	155.6	0.077
	2.50	0.000	0.000	191.1	0.004	0.948	0.000	4240.0	0.923	0.000	0.024	182.5	0.008
	3.00	0.000	0.000	222.2	0.000	0.984	0.000	3460.0	0.750	0.000	0.008	211.0	0.000
	3.50	0.000	0.000	249.4	0.000	0.996	0.000	NaN	NaN	0.000	0.000	235.1	0.000
2	2.00	0.000	0.008	253.2	0.008	0.844	0.000	5673.3	0.923	0.000	0.024	235.9	0.045
	2.50	0.000	0.004	301.0	0.000	0.964	0.000	640.0	0.889	0.000	0.000	280.5	0.008
	3.00	0.000	0.000	354.9	0.000	0.980	0.000	3600.0	0.800	0.000	0.000	334.2	0.000
	3.50	0.000	0.000	395.0	0.000	0.996	0.000	NaN	NaN	0.000	0.000	371.8	0.000

TABLE II
EXPERIMENTAL RESULTS ON SCENARIO 2.

by our CDS. Second, the AR solution suffers from a very high number of FN due to the peculiar structure of the predictive model that encompasses only data coming from a single sensor. Here, the autoregressive component of the model might shadow the additive component of the contaminant source leading to a large amount of false negatives. Differently, the X solution achieves a performance that is similar to the proposed solution even if it is characterized by a larger FPR , that are probably due to the inability to effectively model the datastreams (the autoregressive component here is not considered). Third, the isolation error ϵ_{iso} of the proposed solution is close to zero meaning that in almost all of the cases, the contaminant source is correctly isolated. Differently, the AR solution does not provide satisfactory results, while the X one guarantees results slightly worse than the proposed CDS. Forth, as expected, increasing the parameter Γ of the ICI-based CDT results in a decrease of the FPR at the expense of an increase of FNR and DD for all the solutions. Finally, as expected, FNR , FPR , DD and ϵ_{iso} increase with σ for all the solutions.

Figure 6 compares the distributions of the contaminant measurements in the source zone at the detection time \hat{T} (the blue solid line) and in the training set TS (the red dotted line). Here, we considered only runs where the contaminant was correctly detected and isolated in the runs over *Scenario 1*

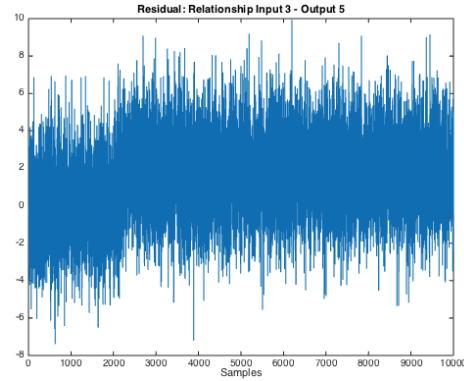


Fig. 5. An example of computed residual for the relationship (*Input Z3-Output Z5*) in Scenario 1 with $\sigma = 2$.

dataset with $\Gamma = 2.5$ and σ ranging from 0.5 to 2. The overlap between the two distributions indicates that a threshold-based approach would not be very effective, as this would induce either many false positive detections or many false negatives. As expected, the overlapping area between the two distribution increases with σ , meaning that contaminant detection becomes more challenging at when noise increases (Figure 6 (c)).

V. CONCLUSIONS

We present a novel CDS that, differently from what is currently present in the literature, addresses the challenging problem of detecting contaminant sources inside a smart building even when the contaminant of interest is also naturally present in the building environment (e.g., the CO₂). Our experiments show that the proposed system can effectively detect and isolate contaminant sources, and that, to this purpose, it is essential to exploit both the temporal and the spatial correlations characterizing the streams of measurements in different zones of the building. Remarkably, the proposed CDS does not require a detailed model of the contaminant/airflows propagation inside the building.

In the future, we plan to extend the proposed methodology to the detection and isolation of sensor faults within smart buildings, which is a particularly challenging scenario that involves the simultaneous presence of both sensor faults and contaminant sources. In addition, we are also considering cognitive reasoning on the affected relationships of the propagation graph in order to differentiate between the two cases.

REFERENCES

- [1] D. Cousins and S. D. Campbell, "Protecting buildings against airborne contamination," *Lincoln Laboratory Journal*, vol. 17, no. 1, 2007.
- [2] T. H. Jeys, W. D. Herzog, J. D. Hybl, R. N. Czerwinski, and A. Sanchez, "Advanced trigger development," *Lincoln Laboratory Journal*, vol. 17, no. 1, pp. 29–62, 2007.
- [3] X. Liu and Z. Zhai, "Prompt tracking of indoor airborne contaminant source location with probability-based inverse multi-zone modeling," *Building and Environment*, vol. 44, no. 6, pp. 1135–1143, 2009.
- [4] M. Michaelides, V. Reppa, M. Christodoulou, C. Panayiotou, and M. Polycarpou, "Contaminant event monitoring in multi-zone buildings using the state-space method," *Building and Environment*, vol. 71, no. 0, pp. 140–152, 2014.
- [5] M. Sohn, R. Sextro, A. Gadgil, and J. Daisey, "Responding to sudden pollutant releases in office buildings: 1. Framework and analysis tools," *Indoor Air*, vol. 13, no. 3, pp. 267–276, 2003.
- [6] G. Boracchi, M. Michaelides, and M. Roveri, "A cognitive monitoring system for contaminant detection in intelligent buildings," in *International Joint Conference on Neural Networks (IJCNN 2014)*, 2014.
- [7] L. Ljung, *System identification*. Springer, 1998.
- [8] C. Alippi, G. Boracchi, and M. Roveri, "A Just-In-Time adaptive classification system based on the Intersection of Confidence Intervals rule," *Neural Networks*, vol. 24, no. 8, pp. 791 – 800, 2011.
- [9] A. Goldenshluger and A. Nemirovski, "On spatial adaptive estimation of nonparametric regression," *Math. Meth. Statistics*, vol. 6, pp. 135–170, 1997.
- [10] C. Alippi, G. Boracchi, and M. Roveri, "Ensembles of change-point methods to estimate the change point in residual sequences," *Soft Computing*, vol. 17, no. 11, pp. 1971–1981, 2013. [Online]. Available: <http://dx.doi.org/10.1007/s00500-013-1130-7>
- [11] R. A. Johnson and D. W. Wichern, *Applied multivariate statistical analysis*. Prentice hall Englewood Cliffs, NJ, 1992, vol. 4.
- [12] M. P. Michaelides, D. G. Eliades, M. Christodoulou, M. Kyriakou, C. Panayiotou, and M. Polycarpou, "A Matlab-CONTAM toolbox for contaminant event monitoring in intelligent buildings," in *Artificial Intelligence Applications and Innovations*. Springer, 2013, pp. 605–614.
- [13] L. Wang, W. Dols, and Q. Chen, "Using CFD capabilities of CONTAM 3.0 for simulating airflow and contaminant transport in and around buildings," *HVAC&R Research*, vol. 16, no. 6, pp. 749–763, 2010.

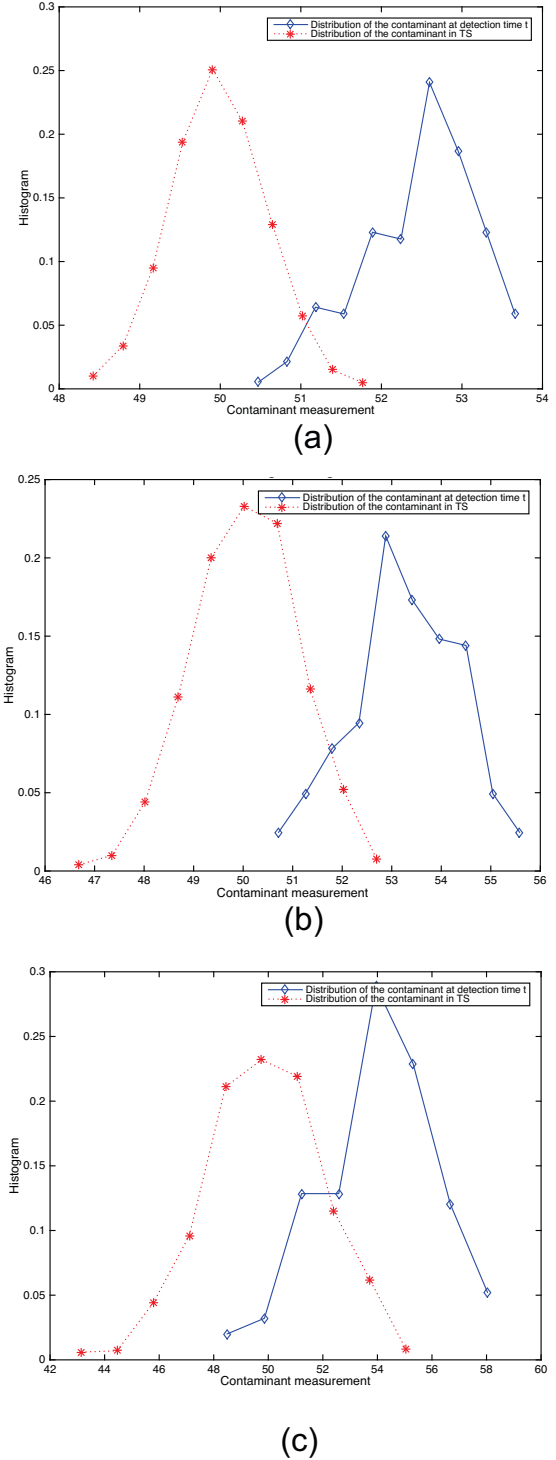


Fig. 6. Distribution of the contaminant measurement at detection time \hat{T} for all the experiments where the contaminant has been correctly detected and isolated (blue solid line) and distribution of the contaminant measurements up to time $t_0 = 1000$ (i.e., the training set) for the zone where the contaminant source has been inserted (i.e., the source zone). The considered scenario is number 1 and $\Gamma = 2.5$. Figure (a), (b) and (c) refer to $\sigma = 0.5$, $\sigma = 1$, and $\sigma = 2$, respectively.