

SPARK: Smart Building Fire Prediction And Risk Anal^ysis

Anik Pramanik*, Mukesh Kumar[†], Nisha Panwar[‡], Laramie V. Potts*, Jainam Shah*, Shantanu Sharma*

*New Jersey Institute of Technology, Newark, NJ, USA

[†]University of California, Merced, CA, USA

[‡]Augusta University, Augusta, GA, USA

Abstract—This paper develops SPARK — a novel algorithm for predicting fire spread in smart buildings. SPARK leverages prior information about the building (e.g., room flammability, structural layout, and occupant distribution) together with real-time multimodal sensing data from static sensors (e.g., temperature sensors, smoke detectors, and embedded structural monitoring devices), dynamic sensors (e.g., drones and robots), and human sensors such as firefighters. By continuously integrating these heterogeneous information streams, SPARK produces a real-time and adaptive prediction of fire spread. Specifically, SPARK classifies rooms into risk zones (e.g., safe, vulnerable, and critical) using a clustering algorithm.

Index Terms—Smart buildings, Sensor data, graph model

1. INTRODUCTION

Fires in high-rise and large-scale buildings remain one of the most severe threats to human life and urban infrastructure. Historical events such as the September 11, 2001, terrorist attacks in New York City [1], the Dubai high-rise fire in 2025 [2], and Grenfell Tower fire in London [3] illustrate the devastating scale of such incidents. Collectively, these disasters accounted for more than 3,000 fatalities, tens of thousands of injuries, and billions of dollars in property and economic losses [4]. These events demonstrate how quickly localized ignition can escalate into large-scale disasters, due to several interconnected factors, including smoke propagation through ventilation systems, the weakening of structural components under extreme heat, and the obstruction of evacuation routes. For instance, in the case of the 9/11 attacks, many firefighters and building occupants died due to the lack of structural information and improper communication on radios; see Chapter 8 of [4]. Similarly, during the Grenfell Tower fire, the persistence of the “stay put” strategy, due to the lack of situational awareness, contributed significantly to the death of 72 people [5].

In modern times, high-rise buildings are increasingly equipped with a variety of sensors and smart devices, including smoke and heat detectors, surveillance cameras, motion sensors, and WiFi access points. These sensors enable real-time monitoring of fire spread, structural integrity, and occupant locations, resulting in an advanced system compared to the conventional fire alarms and passive safety designs. By providing accurate, real-time situational information, smart buildings allow occupants and firefighters to adapt their evacuation and rescue strategies more effectively.

Several works have explored technology-driven fire detection and monitoring in buildings. To evaluate the limitations of these approaches, it is important to first understand the types of sensors available in a smart building. In general, sensors can be classified into three categories: (i) **Static sensors**: are permanently installed in the building and continuously monitor environmental conditions — examples include smoke detectors, heat/temperature sensors, gas sensors, Wi-Fi access points, security cameras, and motion detectors; (ii) **Dynamic sensors**: are mobile devices that can be deployed on demand to gather information in areas where static sensors are unavailable, damaged, or compromised by fire — examples include aerial drones, robotic platforms (e.g., robotic dogs); (iii) **Human sensors**: consist of building occupants or firefighters, who can provide critical real-time information about room conditions, fire spread, or trapped individuals.

While previous works [6]–[12] have contributed to smart building fire prediction and monitoring, the existing smart building fire prediction methods often ignore accurate, location-aware 3D models of buildings, campuses, and cities, created by licensed geomatics professionals, and suffer from over-reliance on isolated sensor modalities, insufficient fusion of static, dynamic, and human sensor inputs, and inadequate modeling of complex building topologies, particularly in cases of indirect fire spread through airflow pathways. For instance, [6] focuses on enhancing situational awareness; however, faces challenges in real-time comprehensive data fusion across multiple heterogeneous sensors. [7] proposes a hybrid ensemble machine learning model using multimodal sensor data, but acknowledges limitations in dynamic adaptability. [12] focuses on fire detection based only on video recorders. [8] designs a multi-information fusion system for intelligent electrical fire detection in green buildings, however, underrepresents the integration of occupant behavior and spatial fire propagation modeling. [9] develops a drone-assisted high-rise monitoring framework to improve fire safety. [10] introduces an intelligent multi-sensor detection system for indoor fires, but faces challenges regarding sensor fusion accuracy and latency in complex fire dynamics. [11] designs an IoT data exchange middleware for emergency response prioritization; however, it lacks integration with adaptive fire risk modeling, which is vital for informed decision-making.

In short, the existing work has limitations in terms of low accuracy and robustness in real-time fire risk assessment, as it relies solely on limited information from either sensors or the building itself. This paper addresses such limitations and provides an answer to the following question:

Can we develop a system for fire prediction that accounts for multiple types of sensors as well as building infrastructure?

We develop SPARK that integrates data from heterogeneous sensors—including static, dynamic, and human sensors—to accurately predict fire spread in smart buildings. Beyond sensor data, SPARK incorporates room attributes such as occupant density, material flammability, ventilation pathways (e.g., air ducts and HVAC systems), and gas conduits. These multimodal inputs are normalized and mapped over a *building connectivity graph*, where nodes correspond to rooms and edges capture both physical connections (e.g., doors, walls, and corridors) and virtual linkages (e.g., airflow ducts and shared utility channels). A diffusion-based propagation model is then applied on this graph to estimate the dynamic spread of fire and smoke, enabling risk levels in one location to influence projections in spatially or virtually adjacent areas. The outcome of this process is the formation of risk-aware room clusters that highlight rooms already compromised as well as those at high probability of imminent spread. By transforming raw sensor data into actionable intelligence, SPARK facilitates a prioritized, adaptive response strategy that supports occupant evacuation, firefighter decision-making, and optimal allocation of emergency resources in real-time.

2. SETTINGS

The section describes the setting of our SPARK algorithm by identifying the available data sources, clarifying the architectural connectivity of the building, and outlining other key critical fire-risk elements. These assumptions and inputs provide the foundation on which the SPARK algorithm is developed, as detailed in the following section.

Data Sources. We consider a high-rise smart building scenario with N rooms for our SPARK algorithm. Each room in such buildings is typically equipped with multiple sensors that support both everyday operations and safety-critical monitoring. For example, WiFi routers not only provide Internet connectivity but can also enable localization and occupancy detection; motion detectors track human movement across rooms; and temperature sensors monitor heat variations that may indicate abnormal activity. In addition, several fire-related sensors play a crucial role in early detection. Smoke/PM sensors measure particle density in the air to identify smoke buildup, CO and CO₂ detectors monitor toxic gas concentrations, and flame/IR sensors detect direct flame signatures. These devices are permanently installed and therefore fall into the category of *static sensors*. During a fire, firefighters may be deployed to directly assess conditions within individual rooms. Their on-scene input, based on situational awareness and professional

expertise, constitutes the category of *human assessments*. Furthermore, mobile platforms such as drones or robotic units can be deployed to collect additional information in real time. These dynamic sensors can detect smoke thickness, fire intensity, human movement, and environmental obstacles, even in areas where static sensors may have failed or provide limited coverage. Such sources are categorized as *dynamic/mobile sensors*, providing complementary observations that enhance the overall fire status assessment.

Critical Rooms. We also account for additional risk factors that increase the likelihood of fire spread or demand heightened attention. Certain rooms present elevated fire risks due to their contents, which can be quantified by a flammability or fuel-load score (e.g., storage areas containing chemicals or archives filled with paper stacks). Other rooms are deemed critical, such as control rooms essential for building operations or spaces where large numbers of people may become trapped by smoke. We assume that we maintain critical room information in advance.

Room Connectivity. To accurately characterize fire propagation within a building, it is essential to account for how rooms are interconnected, as fire spreads from one room to the next. The spatial layout of a building introduces significant complexity because fire and smoke can spread through multiple pathways. On the one hand, rooms may be directly connected by physical adjacency, such as shared walls, doors, or corridors. On the other hand, buildings often contain ventilation systems, such as HVAC ducts, which act as indirect but highly influential channels for fire and smoke spread.

3. SPARK ALGORITHM

A. High-level overview of the Algorithm

Our proposed algorithm begins by constructing the building connectivity structure (§B.1), which combines physical adjacency relations (i.e., walls, doors, corridors) with indirect pathways (e.g., HVAC ducts). Figure 1 gives the overview of our SPARK algorithm. The integrated topology of SPARK forms the basis of a graph representation where nodes correspond to rooms and edges encode potential fire transmission pathways. Then, SPARK algorithm (see Figure 1) for predicting the fire spread in a smart building works as follows:

- 1) **Data collection (§C.1).** Gather readings from static and dynamic sensors (e.g., temperature, smoke, CO₂, drone, firefighters/humans) to assess room-level fire status.
- 2) **Sensor Data Normalization (§C.1).** Apply min–max normalization to map sensor data to a comparable $[0, 1]$ scale.
- 3) **Sensor Data Fusion (§C.2).** Aggregate normalized readings using weighted fusion to ensure robustness in fire assessment even with missing sensors.
- 4) **Fire Propagation with respect to room attributes (§C.3).** Combine fused sensor data with room-level attributes (flammability and criticality) to appropriately scale risk.




Fig. 1. Overview of the SPARK algorithm, where the first block denotes how the buildings are represented and the second block how the fire risk of each room is assessed to detect the rooms that are most at critical condition.

- 5) **Recursive Fire Propagation (§C.4).** Incorporate multi-hop neighbor influence via the graph diffusion to ensure that fire risk accounts for direct and indirect spread.
- 6) **Clustering (§C.5).** Apply k -mean clustering algorithm to partition rooms into *safe*, *vulnerable*, and *critical* clusters for actionable decision support.

B. Modeling a Building

To formalize our algorithm, we develop a mathematical framework that models the building structure, heterogeneous data sources, and fire dynamics, as detailed below:

B.1 Graph Representation of Buildings.

Let N be the number of rooms in a building. To capture the main pathways through which fire and smoke propagate, we focus on two dominant types of adjacency information: structural adjacency and airflow adjacency. These are chosen because they represent the primary mechanisms of spread—flames and heat transfer through shared physical boundaries, and smoke or hot gases traveling through ventilation systems. Other possible connections (e.g., plumbing shafts or electrical conduits) are generally less influential and can be subsumed under these two categories described below:

- **Structure adjacency matrix.** This matrix represents the physical adjacency of rooms, i.e., rooms sharing walls, doors, corridors, or windows.

$$A^{(structure)} \in \{0, 1\}^{N \times N},$$

where $A_{ij}^{(structure)} = 1$ if rooms i and j share a direct physical connection; otherwise 0.

- **Airflow adjacency matrix.** This matrix captures airflow-based adjacency, i.e., connectivity between rooms through ventilation ducts, HVAC systems, or similar channels.

$$A^{(airflow)} \in \{0, 1\}^{N \times N},$$

where $A_{ij}^{(airflow)} = 1$ if rooms i and j are connected through ventilation ducts or HVAC systems, and 0 otherwise.

To create a single and coherent framework for modeling fire spread, we unify the two forms of connectivity (structural and airflow) using a *union rule*: if two rooms are connected through either pathway, they are treated as adjacent in the unified graph. Formally,

$$Adj[*] = \min(A^{(structure)} + A^{(airflow)}, 1). \quad (1)$$

This formulation ensures that $Adj[i, j] = 1$ whenever a connection exists in either the structural or airflow matrix, without prioritizing one type of pathway over the other.

Formally, the building is represented as a graph $G = (V, E)$, where V is the set of rooms and E is the set of edges determined by the nonzero entries of the aggregated adjacency matrix $Adj[*]$. Here, $Adj[i, j]$ denotes the edge between two rooms i and j , capturing both direct and indirect fire-propagation routes.

B.2 Room Attributes.

Suppose the given building consists of N rooms, indexed by $i \in \{1, \dots, N\}$. Each room i is associated with two attributes that capture both its physical vulnerability to fire and its operational importance, denoted as follows:

$$Fuel = (Fuel_1, \dots, Fuel_N) \quad (2)$$

$$Critical = (Critical_1, \dots, Critical_N) \quad (3)$$

- **$Fuel_i \in [1, \infty)$: Fuel load or flammability.** This parameter quantifies the amount and type of combustible material contained in a room. For instance, rooms used for storage may contain highly flammable chemicals, stacks of paper, or plastics, all of which increase fire susceptibility. A higher value of $Fuel_i$ reflects a greater probability of ignition given exposure to fire or high heat. In practice, this attribute can be estimated from the building's inventory of

materials stored in each room. The National Fire Protection Association (NFPA) provides guidelines and reference data that rank materials by their relative flammability. We use this ranking to assign relative weights to the materials present in each room, ensuring that highly combustible substances (e.g., gasoline) are given greater influence than less flammable ones (e.g., paper). This approach prevents underestimating fire risk by treating all items equally and produces a more realistic representation of room-level fuel load. The weighted flammability values of all materials in a room are then aggregated to compute the final $Fuel_i$ score for that room.

- **Criticality** $i \in [1, \infty)$: **Room criticality**. This parameter characterizes the functional or human importance of a room within the building. Criticality does not independently cause fire risk; rather, it modulates the priority of intervention when the room is already threatened by fire or smoke.

Criticality may arise from two perspectives. First, there is *operational importance*, where the room contains essential infrastructure such as control centers, power distribution hubs, or communication facilities whose loss would severely impact building operations. Second, there is *occupant safety*, where the room is frequently occupied by a large number of people or is a known location where certain individuals may be trapped, requiring rapid intervention. In certain environments, the criticality of a room may be determined by the vulnerability of its occupants. For example, in hospitals, patients may be unable to evacuate during a fire alarm; in schools, children may struggle to evacuate safely without assistance; and in other cases, individuals may be trapped inside a room due to fire, smoke, or physical obstacles. To assign an appropriate value to $Critical_i$ for each room i , we consider both pre-existing knowledge and dynamic assessments. Pre-existing knowledge may include occupancy information about vulnerable populations (e.g., hospital wards or classrooms). In addition, dynamic mechanisms such as Wi-Fi localization [13], [14], cellphone signal tracking [15], and camera-based monitoring [16] can be used to estimate the number of individuals currently trapped. Based on this occupancy estimate, a criticality score is assigned to each room, where larger values of $Critical_i$ indicate greater urgency for intervention.

C. Details of the Algorithm

This section provides details of the above-mentioned steps.

C.1 Sensor Data Collection and Normalization.

We assume that evidence of fire in each room is obtained from three complementary categories of observations: static sensors, dynamic sensors, and human assessments. The details of these data sources, along with how their sensor readings are collected and processed, are described below:

- **Static Sensors**. Let $\mathcal{M}_s$ denote the set of permanently installed sensor modalities (e.g., smoke/PM, temperature, CO/CO₂, flame/IR, sprinkler flow). In practice, different modalities measure different quantities in heterogeneous scales and units—for example, temperature in degrees Celsius (°C) and smoke concentration in parts per million (ppm). Because of this heterogeneity in both units and reporting behavior, normalization is necessary to make readings from different modalities comparable. For modeling purposes, we define an *epoch* (e.g., 2 seconds) as the interval between two ticks, denoted as epoch t . Data are collected for each epoch ($t = 1, 2, 3, \dots$) and if sensors operate at different reporting frequencies, we align them by taking the most recent reading available at the end of each epoch. For each modality $m \in \mathcal{M}_s$ (e.g., smoke sensor, temperature sensor), let $r_{i,m}(t)$ denote its raw reading for room i at time t , where $r_{i,m}(t)$ is a scalar value corresponding to that sensor's output at epoch t . If a room contains multiple sensors of the same modality, we take the maximum reading to produce a single value per modality. Each modality reading is then mapped to a common $[0, 1]$ score via a calibration function:

$$static_{i,m}(t) = \mathcal{N}_m(r_{i,m}(t)) \in [0, 1]. \quad (4)$$

In Eq 4, $\mathcal{N}_m(\cdot)$ is a modality-specific normalization function that maps raw readings into the unit interval $[0, 1]$. Each modality m has its own calibration curve, reflecting sensor thresholds, baselines, and operating ranges. For example, $\mathcal{N}_{temp}$ may linearly map temperatures between a baseline of 50°C (no significant risk) and a threshold of 200°C (high risk) to $[0, 1]$, while $\mathcal{N}_{smoke}$ may map smoke concentrations between a baseline of 0 ppm and an alarm threshold of 300 ppm to $[0, 1]$. This baseline–threshold calibration ensures that signals from heterogeneous modalities are placed on a comparable scale. We next combine the modality-specific calibrated signals within each room to compute a unified fire risk assessment. Since not all rooms are equipped with the same set of sensors, the aggregation is performed only over the modalities that are actually present in each room. We assign nonnegative modality weights $w_m^{(static)}$ to each sensor modality, based on factors such as responsiveness, ability to distinguish true fire events from nuisance signals, and robustness under heat or smoke stress. For example, flame/IR and smoke/PM sensors receive higher weights due to their rapid and specific response, while temperature or rate-of-rise sensors are weighted moderately because they react more slowly. CO/CO₂ and sprinkler-flow contribute useful but sometimes lagged signals, whereas humidity or acoustic sensors carry lower weights. The weights are normalized so that, within each room, they always add up to one. Taking all these factors into account, the final fused sensor reading for room i is computed as:

$$static_i(t) = \frac{\sum_{m \in \mathcal{M}_s(i)} w_m^{(static)} static_{i,m}(t)}{\sum_{m \in \mathcal{M}_s(i)} w_m^{(static)}}, \quad (5)$$

where $\mathcal{M}_s(i) \subseteq \mathcal{M}_s$ denotes the set of static modalities available in room i . Intuitively, Eq. 5 computes the weighted aggregation of available sensor readings for the room i .

- **Dynamic or Mobile Sensors.** Let $\mathcal{M}_d$ denote the set of mobile sensor modalities deployed on platforms such as drones, quadrupeds, or other robotic units. These systems provide real-time information in areas where static sensors are absent or have failed (e.g., burned out or disconnected). Typical modalities include thermal imaging, RGB flame segmentation, smoke density estimation, gas concentration sensing, acoustic cues, as well as image and video streams. Each mobile reading is time-stamped and paired with platform coordinates (2D/3D position and floor index), which are then used to assign the reading to a specific room. The platform coordinates are first projected into the building's reference frame, where each room is mapped using polygonal boundaries and floor height ranges. If the mobile device's position falls within a room polygon, the reading is assigned to that room; if it lies near a boundary, it is assigned to the nearest adjacent room; and if it cannot be matched to any room, the reading is discarded. This provides a simple and robust mapping from mobile coordinates to rooms.

Next, similar to the static case (Eq. 4), raw heterogeneous readings from each mobile sensor $k \in \mathcal{M}_d$ are normalized to $[0, 1]$ using the modality-specific function $\mathcal{N}_k(\cdot)$. Then, as in Eq. 5, the normalized modalities within each room are assigned weights $w_k^{(dynamic)}$, reflecting their responsiveness, reliability, and ability to reject nuisance signals. The weighted values are aggregated to yield a representative dynamic score for each room:

$$dynamic_i(t) = \frac{\sum_{k \in \mathcal{M}_d} w_k^{(dynamic)} dynamic_{i,k}(t)}{\sum_{k \in \mathcal{M}_d} w_k^{(dynamic)}}, \quad (6)$$

where $dynamic_{i,k}(t)$ is the normalized reading of modality k in room i , and $dynamic_i(t) \in [0, 1]$ is the aggregated score for room i . If no dynamic modalities are available, we set $dynamic_i(t) = 0$. Similar to the static case, Eq. 6 fuses the available mobile sensor signals for room i , assigning greater weight to modalities that provide faster or more reliable indicators of fire (e.g., thermal imaging over acoustic cues).

- **Human Sensors or Firefighter Assessments.** During a fire, when firefighters are deployed into the building, they can provide direct assessments of fire conditions within a room. These assessments are expressed on a normalized $[0, 1]$ scale, for example: none (0), light smoke (0.25), localized flame (0.5), room involvement (0.75), and full involvement (1.0). The assessment is denoted as $firefighter_{i,j}(t) \in [0, 1]$, where j indexes individual firefighters present in room i at time t . If multiple firefighters

provide input for the same room, their assessments are combined by simple averaging:

$$human_i(t) = \frac{1}{M_i(t)} \sum_{j=1}^{M_i(t)} firefighter_{i,j}(t), \quad (7)$$

where $M_i(t)$ is the number of firefighters reporting from room i at time t . This produces a single human-derived fire estimate $human_i(t) \in [0, 1]$ per room, which is then incorporated alongside static and dynamic sensor aggregates in the evidence fusion stage.

C.2 Sensor Data Fusion.

To obtain a unified picture of fire conditions in each room, we combine the three categories of available information: static sensors, dynamic/mobile sensors, and human (firefighter) assessments. Since, in real settings, some sources provide more reliable or decisive information than others, we assign weights to balance their contributions. For example, firefighter observations carry more weight due to their direct situational awareness, static sensors provide stable continuous monitoring, and mobile platforms such as drones add valuable but sometimes intermittent coverage. At each time step, these already normalized category-level scores are fused into a single room-level estimate:

$$e_i(t) = \alpha_s static_i(t) + \alpha_d dynamic_i(t) + \alpha_h human_i(t). \quad (8)$$

Here, the weights are normalized so they satisfy the constraint $\alpha_s + \alpha_d + \alpha_h = 1$. The weights $(\alpha_s, \alpha_d, \alpha_h)$ are automatically adjusted depending on which sources are present. This guarantees consistency even when some sources are missing. This way, we compute each room's condition against fire, based only on sensor readings, denoted as $E(t) = [e_1(t), \dots, e_N(t)]$.

C.3 Fire Propagation.

While Eq. 8 depicts the condition of a room i , it does not capture two things: first, the impact of criticality and flammability on the room i , and second, the impact of the neighbor rooms. To deal with these situations, SPARK works as follows:

$$s_i(t) = e_i(t) (\phi Fuel_i + \psi Critical_i), \quad (9)$$

where $\phi, \psi \geq 0$ are coefficients that control the relative influence of fuel load and room criticality. The vector $S(t) = [s_1(t), \dots, s_N(t)]$ represents the scaled sensor, adjusted for room-level attributes. However, this scaled sensor score alone does not fully capture the fire risk of a room, due to not incorporating the fact that fire may still spread into the room from neighbors through structural connections (e.g., air ducts). Thus, to properly estimate the fire risk of a room, SPARK accounts not only for its own sensor score of the room but also for the scores of its neighboring rooms. Incorporating neighbor influence, the fire risk $fire'_i(t)$ of room i at time t is defined as:

$$fire'_i(t) = s_i(t) + \sum_{j \in N} Adj[i, j] \times s_j(t), \quad (10)$$

where $Adj[*,*]$ is the adjacency matrix encoding connectivity between rooms (see Eq. 1). In Eq. 10, the fire risk of room i is expressed as the sum of its own scaled sensor score $s_i(t)$ and the aggregate contribution from its directly connected neighbors j , as given by the i -th row of $Adj[*,*]$ multiplied with the vector $S(t)$.

C.4 Recursive Fire Propagation.

The neighbor-augmented risk in Eq. 10 considers only *direct* neighbors, but fire can also spread indirectly through multi-hop connections in the building graph. For instance, if room i is linked to j (1-hop) and j to k (2-hop), then k indirectly influences i 's risk. To model this, we adopt a recursive propagation scheme based on a *row-normalized* adjacency matrix and the principle of *graph diffusion* [17].

Row-normalization prevents highly connected rooms from dominating the propagation. Let D be the degree matrix with $D[i,i] = \sum_j Adj[i,j]$, then the normalized adjacency is

$$\widetilde{Adj} = D^{-1} Adj, \quad (11)$$

so that the weights in each row sum to one. This ensures that each room distributes its total connection strength evenly across its neighbors, preventing rooms with many links from exerting disproportionately large influence in the diffusion process. Using this normalized matrix, the one-hop propagation is:

$$fire_i^{1-hop}(t) = s_i(t) + \sum_j \widetilde{Adj}[i,j] s_j(t), \quad (12)$$

which aggregates risk from a room's own scaled score and its immediate neighbors. Extending this recursively to capture longer-range interactions yields a diffusion process:

$$\begin{aligned} fire_i(t) = s_i(t) + \lambda \sum_j \widetilde{Adj}[i,j] s_j(t) \\ + \lambda^2 \sum_j (\widetilde{Adj}^2)[i,j] s_j(t) + \dots \end{aligned} \quad (13)$$

where $\lambda \in (0,1)$ attenuates the influence of distant rooms. This formulation follows the standard *graph diffusion* model [17], [18], compactly capturing both direct and indirect propagation effects¹. Finally, stacking all rooms gives the overall fire-risk vector:

$$Fire(t) = [fire_1(t), \dots, fire_N(t)], \quad (14)$$

where each entry reflects the total, diffusion-weighted fire risk for that room.

C.5 Risk-Zone Clustering.

Once the total fire risk vector $Fire(t)$ has been computed as in Eq. 14, the next step is to interpret these values in

¹This series is equivalent to the *graph diffusion operator* [17], which aggregates all-hop effects in closed form, can be expressed compactly as:

$$fire_i(t) = \sum_j [(I - \lambda \widetilde{Adj})^{-1}][i,j] s_j(t)$$

terms of actionable decision support for firefighting operations, for example by categorizing rooms into *safe*, *vulnerable*, and *critical* clusters. Such grouping can guide evacuation priorities and resource allocation.

To achieve this, we adopt the well-known *k*-means algorithm [19], which partitions the N rooms into k disjoint clusters by minimizing the variance of fire risk values within each cluster. Intuitively, the objective of *k*-means is to ensure that rooms are assigned to the same cluster have similar levels of fire risk, while rooms in different clusters have different levels of fire risk. Formally, the *k*-means objective is:

$$\min_{\{\mu_1, \dots, \mu_k\}} \sum_{i=1}^N \min_{j \in \{1, \dots, k\}} (fire_i(t) - \mu_j)^2 \quad (15)$$

where μ_j denotes the center (or mean) of cluster j , representing the average fire risk value of all rooms assigned to that cluster. In Eq. 15, the goal is to minimize the intra-cluster distance so that each room is grouped with other rooms of similar fire risk. For interpretability, SPARK defines three risk clusters. Rooms assigned to the cluster with the highest center value μ_j are flagged as *critical zones*, requiring immediate intervention. Rooms in the intermediate cluster correspond to *vulnerable zones* where fire spread is likely, while those in the lowest-risk cluster represent relatively *safe zones*. By combining sensor fusion (Eq. 8), scaling (Eq. 9), propagation (Eq. 14), and clustering (Eq. 15), SPARK provides an end-to-end framework for real-time fire risk prediction and categorization in smart buildings.

4. EXPERIMENTAL EVALUATION

This section will first explain how we simulated a smart building (comprising various types of sensors) under fire, and then present experimental evaluations of SPARK algorithm.

A. Settings and Data Generation

Fire Dynamics Simulator. We use Fire Dynamics Simulator (FDS)², an open-source computational fluid dynamics (CFD) tool developed by National Institute of Standards and Technology (NIST) specifically to simulate fire for our experiment. FDS is widely used for research and safety engineering studies to model realistic fire growth, smoke propagation, and sensor responses in complex geometries.

Building structure. We consider two different building configurations with three and ten stories, each containing 20 rooms per floor. Rooms are of size approximately 4 m × 4 m × 3 m. Each room is constructed by partitioning a floor using combustible interior walls. Doors of size ≈ 1.0 m × 2.1 m are placed in all interior partitions on all floors to allow lateral spread.

Fire ignition and simulation time. The simulation runs for 3600 seconds. Each sensor produces readings at every ≈ 0.5

²<https://pages.nist.gov/fds-smv/>

second. Fire is initiated in the lower-left room of Floor 1 by a floor burner delivering 8000 kW/m^2 over a $3 \text{ m} \times 3 \text{ m}$ patch. No other forced ignitions are used. Subsequent fire spread occurs through wall heating and ignition. Five device types are installed at the center of every room: temperature (T), CO and CO_2 volume fractions, soot mass fraction, and visibility.

Dynamic sensor data generation. FDS framework does not natively support mobile or dynamic sensors (e.g., drones or quadruped robots) that reposition during simulation. Furthermore, to the best of our knowledge, there is no other open-source tool that combines both static and dynamic sensors in the context of fire spread. Thus, we simulate dynamic sensors for collecting temperature and smoke in FDS. To do so, At each time t , the dynamic sensor value for a room is computed as the average of its own temperature and smoke readings with those of its adjacent rooms. Further, to mimic real-world uncertainty, we inject noise by multiplying each averaged value by $(1 + \epsilon)$, where ϵ is $\approx 5\%$. This approach yields realistic, noisy dynamic sensor streams that are consistent with the evolving fire conditions.

Hyperparameter settings. We list all hyperparameters in the order of their corresponding equations in §3 and summarize the reasoning behind each selection.

Modality calibration ($\mathcal{N}_m$ in Eq. 4). Raw readings are mapped to $[0, 1]$ via modality-specific calibration: temperature $60^\circ\text{C} \rightarrow 300^\circ\text{C}$, carbon monoxide $5 \times 10^{-5} \rightarrow 10^{-3}$, carbon dioxide $10^{-3} \rightarrow 5 \times 10^{-2}$, soot $0 \rightarrow 10^{-2}$, and visibility $\mathcal{N}_{\text{vis}}(v) = 1 - \frac{v}{30}$ (clipped to $[0, 1]$). These endpoints cover “normal” versus “dangerous” ranges in our simulation, so all sensors share a common scale without saturating too early.

Static-modality fusion weights ($w_m^{(\text{static})}$ in Eq. 5). For static sensors present in a room, we set $w_{\text{temp}}^{(\text{static})} = 0.20$, $w_{\text{CO}}^{(\text{static})} = 0.20$, $w_{\text{CO}_2}^{(\text{static})} = 0.15$, $w_{\text{soot}}^{(\text{static})} = 0.25$, $w_{\text{vis}}^{(\text{static})} = 0.20$, renormalizing over the available $\mathcal{M}_s(i)$. Soot and visibility react quickly to smoke, so they get slightly higher weight; temperature and CO are strong but can lag; CO_2 is informative but less specific to fire.

Sensor fusion ($\alpha_s, \alpha_d, \alpha_h$) in Eq. 8. We set $(\alpha_s, \alpha_d, \alpha_h) = (0.5, 0.3, 0.2)$ and renormalize if any source is absent, yielding $e_i(t)$ via Eq. 8. Static sensors are always on and more available throughout the building, so they are assigned the highest weight. Human inputs are precise but sporadic; dynamic sensors add useful detail but can be noisy, more expensive, and less available.

Attribute scaling (ϕ, ψ) in Eq. 9. Room attributes *Fuel* and *Critical* are normalized to $[0, 1]$, and we use $(\phi, \psi) = (0.5, 0.5)$ in Eq. 9. To instantiate *Fuel* in simulation, 20% of rooms are designated to contain combustible materials (higher baseline fuel load).

Multi-hop attenuation λ in Eq. 13. We tested several values of λ in $\{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8\}$ and found that $\lambda = 0.4$ achieved the highest accuracy. This value emphasizes higher-order powers of the row-normalized adjacency in Eq. 13.




Fig. 2. Experiment 1: Prediction accuracy of SPARK over time. Accuracy at epoch t is defined as the percentage of rooms correctly predicted from the beginning of the simulation up to the current epoch.

B. Results

In this section, we present experiments to evaluate the performance of our proposed SPARK algorithm. Each experiment highlights a different aspect of predictive capability:

- 1) *Experiment 1: Accuracy over Time* — Evaluates how overall prediction accuracy changes as the fire spreads.
- 2) *Experiment 2: Impact of mixing sensor data* — Compares performance under different sensor settings.

Experiment 1: Accuracy over Time. This experiment evaluates how prediction accuracy evolves as the fire spreads. At each epoch of duration $t = 0.5$ seconds, SPARK predicts which candidate rooms are most at risk of igniting next. At each epoch, we first find $\text{Fire}(t)$ using Eq 14 and then apply clustering over rooms. Then, we discard all those rooms of the cluster, representing burning rooms.³ We then compare this predicted set of rooms against the subset of rooms that actually ignite (reach *burning state*) during the subsequent epochs. So, if x rooms are predicted by SPARK, we then compare them against the next x rooms that reach the *burning state*.

Accuracy at epoch t represents the percentage of rooms correctly predicted by SPARK at that specific time step. Figure 2 illustrates the prediction accuracy for three- and ten-story buildings, with average accuracies of 96.29% and 95.85%, respectively. The lower green line (ten-story) shows a steady, linear rise in accuracy as more rooms ignite over time; SPARK leverages its multi-hop propagation to capture distant signals and refine predictions. In contrast, the blue line (three-story) remains flat for about 750 seconds, as fire initially spreads through direct contact between adjacent rooms—walls and doors—keeping sensor readings stable. Around 750 seconds, heat and smoke begin traveling through HVAC and airflow systems, broadening detection coverage and resulting in higher accuracy.

³According to NIST Fire Dynamics [20], room temperature above 250°C indicates a burning room. We follow the same and consider all rooms having a temperature above 250°C as burning.

TABLE I
EXPERIMENT 2: PREDICTION ACCURACY (%) UNDER DIFFERENT SENSOR SETTINGS FOR THREE-STORY AND TEN-STORY BUILDINGS.

Building	Time (s)	Sensor setting		
		Static-Only	Dynamic-Only	Dynamic+Static
Three-story	300	95.90	93.40	97.20
	600	96.04	93.15	97.20
	900	96.66	93.49	97.64
	1200	97.12	94.77	98.09
Ten-story	300	95.69	91.59	96.72
	600	95.22	91.03	96.84
	900	96.08	91.81	96.95
	1200	95.31	91.67	97.06

Experiment 2: Impact of mixing different sensor data.

We evaluate the contribution of different sensor sources by comparing three settings: (i) *Static-only*: using only three static sensors (CO, CO₂, smoke) among all available, (ii) *Dynamic-only*: using two dynamic sensors (temperature and visibility), and (iii) *Dynamic and Static*: combining all five sensors. This experiment was conducted by restricting the sensor to these subsets, and prediction accuracy at the current time is measured as the percentage of rooms that are correctly predicted by SPARK from the beginning of the simulation up to the current time (same as *Experiment 1*).

Table I summarizes the prediction accuracy of SPARK under different sensor settings for the two types of buildings. The results show that combining dynamic and static sensor data improves accuracy consistently compared to static-only or dynamic-only sensor data.

5. CONCLUSION

We develop SPARK — a fire prediction system for smart buildings based on various types of sensors, including static, dynamic, and human sensors, as well as the conditions of the room, such as the presence of flammable material in the room. Mixing different sensor readings with building conditions resulted in a maximum accuracy of $\approx 96\%$ in identifying rooms vulnerable to fire.

ACKNOWLEDGMENT

This project was partially funded by an NCEES Education award and NSF Grant 2245374.

REFERENCES

- [1] September 11 attacks: <https://tinyurl.com/yc27um8s>.
- [2] Massive fire rips through 67-story: <https://tinyurl.com/37k6vuh6>.
- [3] Grenfell Tower fire: <https://tinyurl.com/2djbnc9w>.
- [4] Final Report on the Collapse of the World Trade Center Towers: <https://tinyurl.com/4t3dy37h>.
- [5] Residents told to 'stay put for too long': <https://tinyurl.com/323wjmfvf>.
- [6] T.-Y. Fan, F. Liu, J.-W. Fang, N. Venkatasubramanian, and C.-H. Hsu, "Enhancing situational awareness with adaptive firefighting drones: leveraging diverse media types and classifiers," in *Proceedings of the 13th ACM Multimedia Systems Conference*, 2022, pp. 279–286.
- [7] S. Jana and S. K. Shome, "Hybrid ensemble based machine learning for smart building fire detection using multi modal sensor data," pp. 473–496, 2023.
- [8] X. Ren, C. Li, X. Ma, F. Chen, H. Wang, A. Sharma, G. S. Gaba, and M. Masud, "Design of multi-information fusion based intelligent electrical fire detection system for green buildings," *Sustainability*, vol. 13, no. 6, p. 3405, 2021.
- [9] F. Liu, T. Fan, C. Grant, C. Hsu, and N. Venkatasubramanian, "Dragonfly: Drone-assisted high-rise monitoring for fire safety," in *40th International Symposium on Reliable Distributed Systems, SRDS 2021, Chicago, IL, USA, September 20-23, 2021*. IEEE, 2021, pp. 331–342. [Online]. Available: <https://doi.org/10.1109/SRDS53918.2021.00040>
- [10] J. Baek, T. J. Alhindi, Y.-S. Jeong, M. K. Jeong, S. Seo, J. Kang, and Y. Heo, "Intelligent multi-sensor detection system for monitoring indoor building fires," *IEEE Sensors Journal*, vol. 21, no. 24, pp. 27982–27992, 2021.
- [11] K. E. Benson, G. Bouloukakis, C. Grant, V. Issarny, S. Mehrotra, I. D. Moscholios, and N. Venkatasubramanian, "FireTex: a prioritized iot data exchange middleware for emergency response," in *Proceedings of the 19th International Middleware Conference, Middleware 2018, Rennes, France, December 10-14, 2018*, P. Ferreira and L. Shriram, Eds. ACM, 2018, pp. 279–292. [Online]. Available: <https://doi.org/10.1145/3274808.3274830>
- [12] S. Saponara, A. Elhanashi, and A. Gagliardi, "Exploiting R-CNN for video smoke/fire sensing in antifire surveillance indoor and outdoor systems for smart cities," in *IEEE International Conference on Smart Computing, SMARTCOMP 2020, Bologna, Italy, September 14-17, 2020*. IEEE, 2020, pp. 392–397. [Online]. Available: <https://doi.org/10.1109/SMARTCOMP50058.2020.00083>
- [13] S. He and S. G. Chan, "Wi-fi fingerprint-based indoor positioning: Recent advances and comparisons," *IEEE Commun. Surv. Tutorials*, vol. 18, no. 1, pp. 466–490, 2016. [Online]. Available: <https://doi.org/10.1109/COMST.2015.2464084>
- [14] P. Bahl and V. N. Padmanabhan, "RADAR: an in-building rf-based user location and tracking system," in *Proceedings IEEE INFOCOM 2000, The Conference on Computer Communications, Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies, Reaching the Promised Land of Communications, Tel Aviv, Israel, March 26-30, 2000*. IEEE Computer Society, 2000, pp. 775–784. [Online]. Available: <https://doi.org/10.1109/INFCOM.2000.832252>
- [15] Y. Tai and T. Yu, "Using smartphones to locate trapped victims in disasters," *Sensors*, vol. 22, no. 19, p. 7502, 2022. [Online]. Available: <https://doi.org/10.3390/s22197502>
- [16] Y. Zhang, D. Zhou, S. Chen, S. Gao, and Y. Ma, "Single-image crowd counting via multi-column convolutional neural network," in *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2016*, pp. 589–597.
- [17] L. Katz, "A new status index derived from sociometric analysis," *Psychometrika*, vol. 18, no. 1, pp. 39–43, 1953.
- [18] P. Bonacich *et al.*, "Eigenvector-like measures of centrality for asymmetric relations," *Social networks*, vol. 23, no. 3, pp. 191–201, 2001.
- [19] J. MacQueen, "Some methods for classification and analysis of multivariate observations," in *Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, 1967, pp. 281–297.
- [20] National Institute of Standards and Technology (NIST), Fire Dynamics: <https://www.nist.gov/el/fire-research-division-73300/firegov-fire-service/fire-dynamics>, Engineering Laboratory, Fire Research Division. Updated July 2, 2024; originally published November 17, 2010.