Vrije Universiteit Amsterdam

Fire Department Amsterdam-Amstelland

Master Thesis

Data-Driven Simulation of Alternative Deployment Models for the Amsterdam Fire Department

Implications for Operational Efficiency and Responsiveness

Author: Wenjie Hu (2743457)

1st supervisor:	Rob van der Mei	(Vrije Universiteit Amsterdam)
daily supervisor:	Guido Legemaate	(Fire Department Amsterdam-Amstelland)
2nd reader:	Sandjai Bhulai	(Vrije Universiteit Amsterdam)

A thesis submitted in fulfillment of the requirements for the Master of Science degree in in Business Analytics at the Vrije Universiteit Amsterdam.

September 4, 2025

Acknowledgements

This thesis was conducted during my internship at the Information Management team of the Fire Department Amsterdam-Amstelland.

First and foremost, I would like to thank my family for their constant support and care in daily life, which enabled me to focus fully on this research.

I am especially grateful to my daily company supervisor, **Guido Legemaate** from Fire Department Amsterdam-Amstelland, for his invaluable guidance throughout the internship. His advice covered not only the research direction and methodology, but also technical details, expert insights, and access to data — all of which were essential to the success of this thesis.

I would also like to thank my first academic supervisor, **Prof. Rob van der Mei** from Vrije Universiteit Amsterdam, for his encouragement, insightful suggestions, and technical guidance. In addition, I thank my second academic supervisor, **Prof. Sandjai Bhulai** from Vrije Universiteit Amsterdam, for dedicating his time to review and assess this thesis.

Finally, I would like to express my gratitude to all colleagues and everyone who provided support or guidance during my internship and this research, both at the Fire Department Amsterdam-Amstelland and beyond.

Management Summary

This study examined ways for the Fire Department Amsterdam-Amstelland (FDAA) to improve operational efficiency through more flexible deployment models. The research focused on alternative vehicle compositions, adaptive dispatch rules, and relocation strategies, evaluated through an enhanced simulation framework. Performance was measured using the **Fraction of On-Time Arrivals (FOTA)**, a single, interpretable metric that reflects both response speed and coverage.

Simulation experiments demonstrated clear benefits from flexible deployment:

- Introducing smaller and faster vehicles increased FOTA by more than 40% across all incidents.
- Adaptive dispatch rules added a further 11–13% improvement in response performance.
- These gains were consistent even under high-demand scenarios, showing that flexible deployment strengthens both efficiency and resilience.
- Relocation algorithms also helped maintain coverage, but their operational use depends on the development of well-designed trigger mechanisms, which requires further research.

Based on these findings, we recommend that FDAA:

- Gradually introduce lighter units and more flexible dispatch rules into practice.
- Continue using simulation-based evaluation as a tool to support planning and decision-making.

List of Acronyms

Acronym	**Meaning**
FDAA	Fire Department Amsterdam-Amstelland
TS	*Tankautospuit* (Pumper)
RV	*Redvoertuig* (Ladder truck)
HV	*Hulpverleningsvoertuig* (Rescue vehicle)
WO	*Waterongevallenvoertuig* (Marine rescue boat)
MOTO	Motorcycle unit
AED	Automated External Defibrillator
CPR	Cardiopulmonary Resuscitation
MCRP	Maximum Coverage Relocation Problem
LBAP	Linear Bottleneck Assignment Problem
SDOM	Station Dispatch Optimization Model
FOTA	Fraction of On-Time Arrivals

Contents

1

Introduction

1.1 Background

The Fire Department Amsterdam-Amstelland (FDAA) is responsible for providing fire protection and emergency response services across a densely populated region in the Netherlands, which includes the city of Amsterdam and surrounding municipalities such as Amstelveen, Aalsmeer, Diemen, Uithoorn, and Ouder-Amstel. The FDAA operates 19 fire stations, each equipped with various emergency vehicles, such as pumpers and ladders.

The FDAA operates within a complex urban environment with over one million residents. The department handles a wide variety of emergencies including structural fires, traffic accidents, storm damage, and medical incidents. Each fire station is assigned a specific service area and may also contribute to broader response neighbourhoods depending on incident location and severity. Resources are centrally dispatched by a call centre, which alerts stations only after determining the required type and scale of response.

Currently, like the majority of safety regions in the Netherlands, the FDAA dispatches resources using a relatively fixed and streamlined operating model. Regardless of the severity or type of incident, whether it is a large fire, a small fire, or even a medical emergency such as a cardiac arrest, the standard approach is usually to dispatch at least a pumper truck with full six-person crew, in specific cases accompanied by one or more specialized trucks.

1.2 Problem

While current deployment strategy ensures operational readiness, it may not be the most time- or resource-efficient solution, especially in high-frequency, low-severity scenarios. Given the increasing demand and limited availability of personnel, there is a growing interest in more adaptive and modular deployment strategies.

For example, in the case of out-of-hospital cardiac arrests (OHCA), every minute of delay can increase the likelihood of fatality by approximately 7%-10% (1). In the Netherlands, the legally permitted ambulance response time is up to 15 minutes, whereas the fire department is currently required to respond within 5–10 minutes for high-priority emergencies, with a maximum of 18 minutes. As a result, the fire department is frequently called upon for urgent medical support. However, under the current system, a 6-person crew is required because manual CPR is exhausting and must be rotated. With an AED, however, a single responder may suffice, making rapid intervention by a motorcycle unit carrying an AED feasible. This would also prevent tying up an entire crew, preserving availability for other incidents.

1.3 Earlier Studies

In recent years, the FDAA has actively collaborated with academic researchers to improve operational efficiency through data-driven decision-making. A series of studies have explored mathematical modelling, machine learning, and optimization techniques to support tactical and strategic planning.

For example:

- Van den Berg et al. (2) developed an integer linear programming (ILP) model to optimize fire station locations and vehicle allocation, with the objective of maximizing population coverage and minimizing response times across different vehicle types.
- Usanov et al. (3) proposed a dynamic vehicle relocation strategy using the Maximum Coverage Relocation Problem (MCRP) and the Linear Bottleneck Assignment Problem (LBAP), allowing the FDAA to maintain service levels during major incidents.
- Legemaate et al. (4) investigated incident forecasting based on weather data, showing that small incidents can be reasonably predicted using Exponential Averaging (EA), while larger incidents follow a more complex stochastic process.

While these studies have made important contributions to long-term planning and incident management, they primarily address system-level optimization in large-scale or strategic scenarios. By contrast, relatively little attention has been given to operational deployment efficiency in high-frequency, low-severity events—such as medical emergencies or small fires—where full-scale deployment may be unnecessary.

In particular, there is still a gap in the development of modular, adaptive dispatch strategies that incorporate real-time resource availability and incident-specific requirements. Addressing this gap is especially important in context of rising incident volumes and personnel shortages. This study aims to build upon previous work by introducing a flexible deployment model designed to improve response to everyday emergencies, while also ensuring sufficient coverage and preparedness for major incidents across the entire service area.

1.4 Research Objectives

The current resource deployment strategy of the FDAA is based on rigid, predefined unit structures—most commonly the standard dispatch of a six-person pumper team (TS6) to almost all incident types. While this approach ensures readiness for major emergencies, it often leads to over-allocation of personnel, particularly for smaller or lower-severity incidents. These inefficiencies may consequently increase response times, limit resource utilization efficiency, and reduce overall system flexibility.

This research examines whether operational efficiency can be significantly enhanced by adapting the current deployment model. In particular, we investigate the potential benefits of introducing more flexible unit compositions—such as two- and four-person pumpers (TS2, TS4) and single-responder motorcycles (MOTO). These configurations have already been implemented in several countries, demonstrating their practical feasibility for minor incidents and time-critical emergencies such as cardiac arrests. Since real-world experimentation is impractical and potentially unsafe, simulation is essential for assessing their effects on response times and system coverage, and for identifying optimal rules for vehicle dispatch and relocation. However, the simulation tool available to the FDAA in 2019 had several important limitations: it could not model flexible team sizes, simulate mixed-resource dispatching, adapt relocation algorithms to new configurations, or capture dynamic incident escalation. Addressing these constraints forms the core of this research.

The main objective of this research is to improve operational efficiency at the FDAA through more flexible and realistic simulation-based planning. Specifically, this project aims to:

- Enhance the fire department's existing simulator to support flexible vehicle and crew configurations, including new types such as TS2, TS4, TS6, MOTO, and composite units like TS6_P assembled from smaller vehicles.
- Implement adaptive dispatching rules that allow partial or minimum viable deployment for incidents like medical emergencies, where rapid single-responder arrival can be life-saving.
- Extend and adapt relocation optimization models such as the MCRP and the LBAP for relocation during major incidents, moving available units from less busy to temporarily depleted stations to maintain area-wide coverage.
- Integrate an expanded dataset (2007–2025), including the newly established GBA fire station, to enable more realistic scenario testing, model calibration, and validation.
- Evaluate the operational impact of various dispatch and relocation policies using simulation experiments under realistic demand conditions.

This simulator-driven approach allows the FDAA to safely test innovative strategies in a controlled environment, making data-driven decisions before implementing any policy changes in the field.

1.5 Research Question

Based on the objectives outlined above, the overarching research question of this thesis is:

> **Main Research Question:** How can an alternative deployment model be designed for the fire department, and what are its implications for operational efficiency and responsiveness in terms of emergency response and deployment capacity?

This main question can be divided into the following sub-questions:

1. What are the limitations of the current dispatch and relocation mechanisms in the fire department's operational model?

2. How can the simulation framework be extended to incorporate additional vehicle types and flexible dispatch rules?

3. How can historical incident and deployment data be used to calibrate and validate these extensions?

4. What is the impact of implementing dynamic relocation algorithms on station coverage during major incidents?

5. How do the proposed changes affect key performance indicators such as response time, resource utilization, and the need for external assistance?

1.6 Challenges and Considerations

While the use of smaller, more flexible units can theoretically increase responsiveness and reduce unnecessary resource commitment, it also introduces new complexities. For instance:

- Smaller units may reduce the availability of full 6-person teams needed for major incidents.
- The training and certification of firefighters for multiple vehicle types must be adapted.
- Dispatching decisions become more complex and require sophisticated algorithms to balance coverage and risk.

Furthermore, adapting existing relocation algorithms (e.g., MCRP + LBAP) to this new paradigm is non-trivial, as they were originally designed for fixed team configurations. These models must be updated to account for flexible team sizes and partial deployments.

1.7 Thesis Structure

The remainder of this thesis is organized as follows:

- **Chapter 2: Background** – Introduces the theoretical and practical background of fire department operations, dispatch strategies, and existing optimization models.
- **Chapter 3: Literature Review** – Reviews simulator development for Dutch emergency services, historical and FDAA optimization models, and examples of international fire department optimization applications.

- **Chapter 4: Data** – Details the datasets used in the study, including incident records, vehicle and station data, and preprocessing steps.
- **Chapter 5: Algorithm** – Describes the modeling of flexible deployment strategies and the adaptation of relocation algorithms.
- **Chapter 6: Simulation** – Presents the simulation framework and the implementation of the algorithm within the simulator.
- **Chapter 7: Experiment** – Reports experimental results under realistic demand scenarios and analyses operational impacts.
- **Chapter 8: Conclusion** – Summarizes the main findings, key insights, and lessons learned from the research.
- **Chapter 9: Discussion and Recommendation** – Provides a detailed discussion of implications, trade-offs, limitations, and recommendations for operational implementation.

2

Background

The FDAA is responsible for fire protection and emergency response across the Amsterdam-Amstelland region. It operates a diverse fleet—including, but not limiting, pumpers, ladder trucks, rescue vehicles, and marine units—to handle various incident types, such as fires, traffic accidents, and water-related emergencies. Staffing consists of both full-time professionals, who are stationed at the fire station 24/7, and part-time volunteers, who are called in when required.

2.1 Fire Stations

The FDAA operates 19 fire stations, strategically distributed across the region to ensure fast and effective response. Each station is responsible for a specific service area (verzorgingsgebied), and together, these areas form a network of response neighborhoods. The number of stations covering each neighborhood depends on factors such as population density, historical incident frequency and risk.

As shown in Figure 2.1, central and northern Amsterdam have a higher concentration of fire stations, reflecting the historically greater incident frequency in these densely populated areas. By contrast, the southern regions—where emergencies occur less often—are served by fewer stations with more limited resources. For example, Aalsmeer and Uithoorn operate as volunteer-based stations on weekends, relying on part-time personnel during those periods.

It is worth noting that the GBA station (Gezamenlijke Brandweer Amsterdam) began operations on July 1, 2020, after the development of the original simulation tool. For accurate modeling, the simulator must therefore be adapted to include this station. In

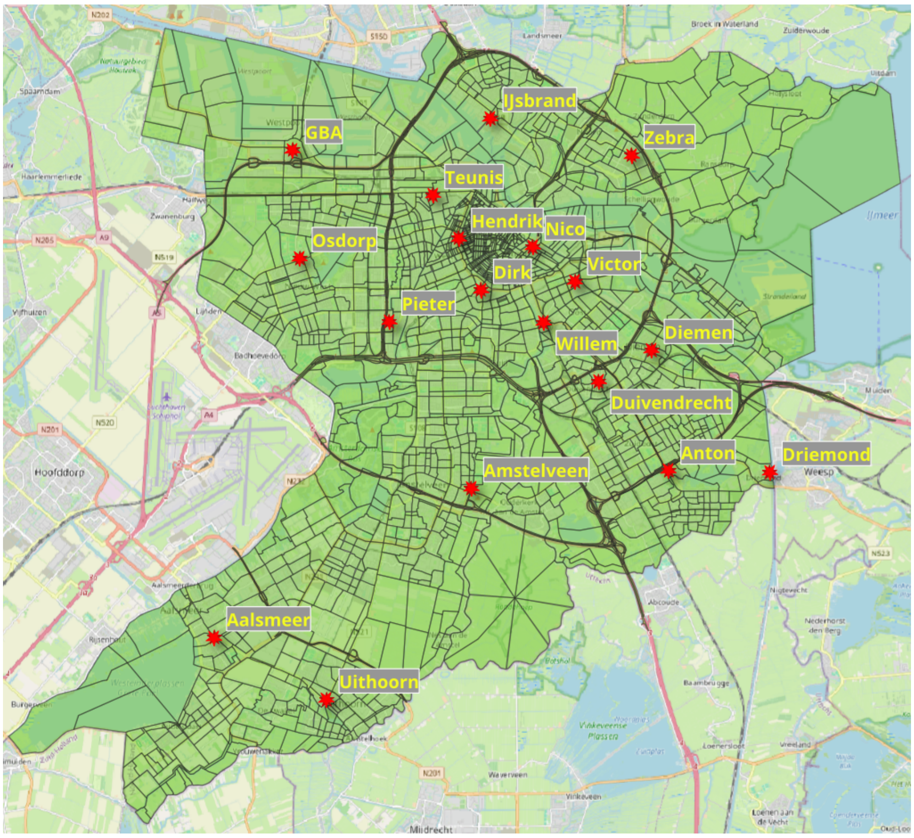

Figure 2.1: Distribution of fire stations in Amsterdam

contrast, the Weesp station, which started operations in 2024, lacks sufficient historical data and is thus excluded from the simulation. Consequently, this station is also not shown in the referenced map (Figure 2.1).

The simulation framework depends on accurate incident and vehicle dispatch records per service area in order to fit proper probabilistic distributions for realistic scenario generation.

2.2 Vehicle Types

The FDAA operates four main types of emergency response vehicles, each designed to handle specific incident types:

- **TS (Tankautospuit / Pumper):** The most frequently dispatched unit, used in most fire-related incidents, typically crewed by six firefighters and fully equipped for general firefighting operations, but can also handle lighter technical rescue operations.

- **RV (Redvoertuig / Ladder truck):** Used for incidents involving high-rise buildings or requiring vertical access and aerial water streams.
- **HV (Hulpverleningsvoertuig / Rescue vehicle):** Equipped for complex technical rescue operations, such as vehicle extrications and roadside accidents, involving heavy equipment.
- **WO (Waterongevallenvoertuig / Marine rescue vehicle):** Deployed to water-related emergencies—such as drownings or boat accidents—and crewed by four specially trained divers.

Each fire station is assigned specific crews and vehicles according to local demand and incident profiles. Staffing levels and vehicle allocations are based on operational requirements and historical incident patterns. For instance, stations in urban centres are more likely to have RVs and HVs due to the greater likelihood of high-rise fires or complex traffic accidents, whereas marine units are positioned near bodies of water to respond to water-related emergencies.

Pumpers are available at every station due to their broad applicability, whereas specialized vehicles such as RV, HV, and WO are stationed only at selected locations. Some incidents require only a single vehicle type, while others demand a combination of units. For example, fires are classified as small, medium, or major, with corresponding deployments of one, two, or three pumpers, respectively.

Each vehicle type requires a specific crew composition, and vehicles can only be dispatched when fully staffed:

- A TS requires 6 personnel: 4 general firefighters, 1 driver and 1 captain.
- Both RV and HV require 2 specialized personnel, referred to as RVHV staff, who are qualified for aerial operations or technical.
- A WO requires 4 certified divers trained in underwater rescue.

These staffing constraints are strictly enforced in the current dispatch system, which means that vehicles are never deployed without a full, qualified crew. While this requirement ensures operational safety, it also poses challenges for maintaining adequate service coverage during peak demand or staff shortages.

2.3 Personnel

In addition to differences in vehicle crew requirements, firefighters in the FDAA are classified into two categories based on their employment status:

- Full-time firefighters are stationed at the fire stations 24/7 and respond to alarms immediately.
- Volunteer (part-time) firefighters live near the fire station and are called in when needed. They must first travel to the station after being alerted, and the vehicle departs only once the required crew size has been assembled.

From a skill standpoint, there is no difference between full-time and volunteer firefighters, as both receive the same professional training. The primary distinction lies in availability at the station, which directly impacts turnout time— the time required for a crew to get ready and depart after being alerted. Vehicles relying on volunteer staffing generally have significantly longer turnout times than those operated by fully staffed full-time units.

For instance, in the case of TS crew composition:

- Driemond and Duivendrecht stations are staffed entirely by volunteers.
- Aalsmeer, Uithoorn, and Diemen operate with full-time crews on weekdays but switch to volunteer staffing on weekends.
- Amstelveen follows a hybrid model, combining both full-time and volunteer firefighters.

In Amstelveen, the dispatch system gives priority to full-time crews, with volunteers mobilized only when additional staffing is required. If the mandated crew size has not been reached, a vehicle must wait for volunteer firefighters to arrive. In our simulator, this rule also applies if relocation temporarily results in a station being staffed by both full-time and volunteer firefighters.

2.4 Incident Types

The FDAA responds to a wide variety of incidents, not limited to fires. These include technical rescues, water-related emergencies, hazardous material situations, medical aid, and more. In the legacy dataset, incident types were recorded using labels such as Binnenbrand (indoor fire), Liftopsluiting (person stuck in elevator), Persoon te water (person in water), Afhijsen spoed (urgent lifting for ambulance), and others. These detailed categories are later merged into a smaller set of generalized incident types (see Data chapter for mapping), enabling more efficient modeling and scenario generation.

Each incident is assigned a priority level, which reflects its urgency. Although our simulator framework primarily uses three levels (1 being the highest), in practice the priority scale is broader, with levels such as 4 and 5, used for less urgent situations. Incidents with a priority level other than 1 are typically classified as non-critical, posing no immediate threat to life or major property loss — such as elevator entrapments or minor trash fires. Only priority 1 incidents are considered urgent, requiring immediate response. For these high-priority cases, legal standards define the maximum allowable response time, depending on the type of building involved. Table 2.1 summarizes these requirements for fire-related incidents:

Table 2.1: Maximum response times (for fire incidents) by building type

Building Type	**Max Response Time**
Enclosed shopping centers, residences above shops, detention facilities	5 minutes
Apartment buildings, housing for persons with reduced self-reliance	6 minutes
Other residential, retail, healthcare, educational, or lodging facilities	8 minutes
Offices, industrial sites, sports/assembly venues, other functions	10 minutes
Justified deviation with regional rationale	18 minutes (max)

To meet these targets, the fire department initiates response procedures immediately after receiving an alert from the dispatch centre. This involves activating the selected vehicle(s) and notifying the required firefighters, who must quickly prepare themselves before departure. This preparation phase, known as the turnout time, includes tasks such as putting on protective clothing, gathering and checking necessary equipment, and boarding the vehicle. The vehicle can depart only once the full crew is ready, marking the end of the turnout time.

For priority 1 incidents, vehicles use sirens and lights and are allowed to take priority lanes and drive at slightly higher than the maximum allowed speed. In contrast, lower-priority incidents do not justify such urgency. Firefighters can take more time to prepare, and vehicles travel under normal traffic regulations. In priority 2 situations, some privileges—such as using emergency or bus lanes—may still apply, but without sirens or lights.

In this study, we primarily focus on priority 1 incidents. However, it is important to note that some records in the dataset may reflect inconsistencies between the initial and recorded priority. For instance, a dispatcher may initially assign a priority 2 status, but after discussion with the field crew and given severe traffic congestion, the crew may be allowed to respond as if it were a priority 1. These cases often result in longer turnout times and overall response times, yet are ultimately logged as priority 1. In our data, such incidents appear as outliers in the response time distribution.

2.5 Response Time

The total response time of the fire department consists of three main components: dispatch time, turnout time, and travel time. Together, these elements represent the timeline from the moment an emergency call is received to the arrival of the fire brigade at the incident location. This process is illustrated in Figures 2.2.

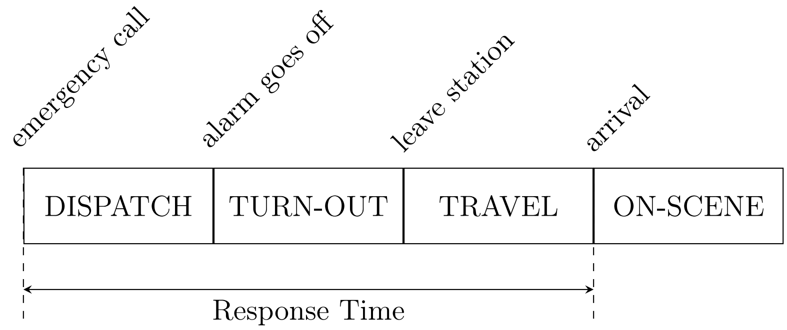

Figure 2.2: Key components of the response time (adapted from (5)).

This decomposition follows the operational timeline of an emergency response, as detailed in *The Data-Driven Fire Department* (5):

- **Dispatch time** refers to the duration between the emergency call being received and the moment when alarms are triggered in the relevant fire station. During this period, the dispatcher collects information from the caller via the emergency dispatch system (GMS), determines the priority level of the incident, and decides which units to deploy.

- **Turnout time** is the interval between the alarm going off in the fire station and the moment the vehicle leaves the station. This period includes firefighter preparation, such as getting dressed, equipping gear, and boarding the vehicle. Once ready, a firefighter presses a button in the vehicle to mark the departure, which is automatically logged in the system.
- **Travel time** begins when the vehicle departs from the station and ends upon arrival at the incident location. The arrival time is again recorded by pressing a designated button in the vehicle.

Together, these three components define the legally and operationally defined response time. These standards are used to ensure sufficient emergency coverage and to evaluate fire department performance across regions.

2.6 Current Operational Model

The FDAA currently operates under a standardized, fixed deployment model. When an incident is reported, the call centre first determines the type of emergency and the required resources, then alerts the relevant fire stations.

Once the alarm is triggered, firefighters immediately prepare—donning protective gear, checking equipment, and boarding the vehicle. Departure occurs only when the full crew is ready, ensuring a structured response. For most routine incidents, a single six-person pumper crew is dispatched, while larger-scale emergencies receive additional vehicles. Minor incidents, including small fires or medical emergencies such as cardiac arrests, typically follow this uniform deployment pattern.

Figure 2.3 illustrates this process, based on internal logs from the operational systems used by Dutch fire services.

2.7 Challenges and Opportunities for Improvement

Despite the clarity and consistency of the FDAA's current dispatch model, efficiency challenges arise when responding to smaller incidents. The default deployment of a six-person

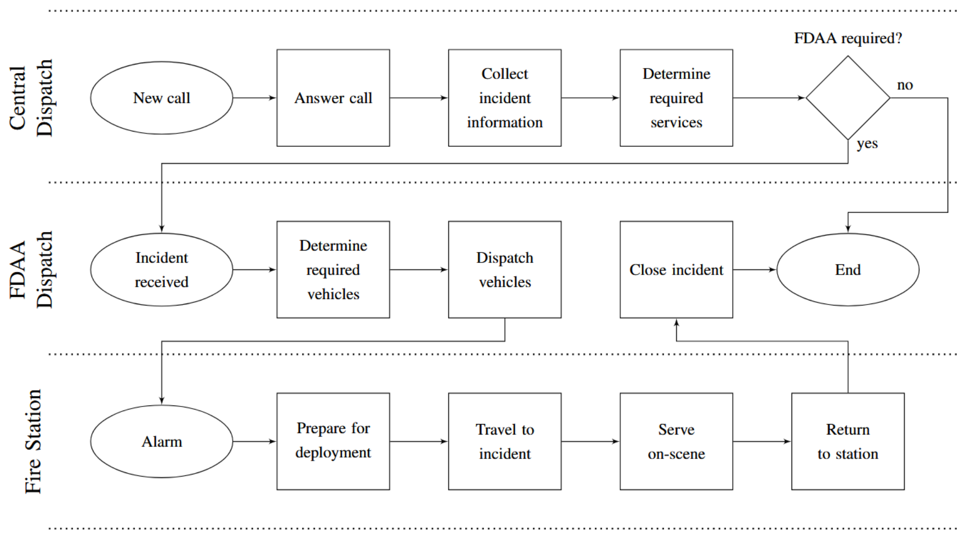

Figure 2.3: Operational timeline of a fire incident response (adapted from (5)).

pumper crew, regardless of incident type or severity, can lead to overuse of resources, particularly for minor incidents where a full crew is not necessary.

One potential solution is the adoption of a more flexible, modular deployment system. Instead of relying solely on standard six-person pumpers, the fire department could deploy smaller units, such as four-person or two-person pumper teams, or even single-respondent motorcycle units, based on the specific needs of each incident.

For instance, in the case of cardiac arrest, medical studies indicate that each minute without defibrillation reduces the patient's survival rate by approximately 7%-10% (1). In the Netherlands, ambulances have a statutory response time of 15 minutes, which runs the risk of arriving too late for effective intervention. By contrast, the fire department's statutory response time is 5–10 minutes for high-priority emergencies, making it a critical first responder.

As discussed earlier, a motorcycle unit carrying an AED can often provide the critical first intervention on its own. Deploying single-person motorcycle units could substantially reduce response time and improve survival outcomes in such time-critical situations. Modular units also help maintain regional coverage. If a full pumper team is engaged in a minor incident, the station may become temporarily unavailable for other emergencies. For example, while responding to a cardiac arrest, the same station might be unable to attend a nearby small fire. By dispatching a motorcycle unit for the medical emergency and a smaller pumper team for the fire, the department can preserve broader operational readiness.

In addition to varying unit sizes, dispatch rules can be refined to allow hybrid responses. In high-priority incidents requiring a full six-person team, nearby stations lacking full personnel or suitable apparatus could dispatch a partial crew to perform initial assessment, evacuation, or communication with the control centre, with the main pumper arriving subsequently. Conversely, if only large pumpers are available, they can deploy a minimal crew—such as sending a firefighter equipped with an AED—rather than waiting for a more distant station with the optimal vehicle. Such adaptive strategies enable faster, more efficient responses while improving overall operational resilience.

2.8 International Inspiration

While six-person pumpers remain standard in Amsterdam, many smaller incidents do not require such a large response. This has prompted research and pilot programs exploring alternative unit configurations. This approach is not purely theoretical. Various fire departments worldwide have already adopted diversified vehicle and staffing models.

In countries such as the United Kingdom and Singapore, fire motorcycles have been successfully deployed to handle time-critical or small-scale incidents. The Merseyside Fire and Rescue Service in the UK launched a pilot scheme using motorcycles equipped with water tanks and hose reels to address anti-social behaviour fires, such as wheelie bins and skips, which represent a significant portion of their call-outs (6). In Singapore, the Civil Defence Force (SCDF) deploys fire bikes manned by Emergency Medical Technicians to rapidly reach life-threatening cases, such as cardiac arrests, ahead of ambulances. These bikes carry essential medical equipment and enable faster initial response in dense urban environments (7).

Another practical example comes from the Lafayette County Fire Department in the United States, which introduced a mini pumper to reduce wear-and-tear on larger vehicles while ensuring their availability for major incidents (8).

These international practices demonstrate the feasibility of diversifying deployment models. By assigning units based on incident type, fire services can allocate resources more effectively, reduce response times, and maximize the impact of available personnel.

2.9 Remaining Challenges

Implementing such changes poses several challenges. New team configurations require appropriate training, and in some cases, new vehicles and equipment. Moreover, splitting teams for small incidents may temporarily limit a station's capacity to respond to larger fires. Although rare, high-impact cases could still lead to delays in responding to severe incidents, especially if nearby stations are simultaneously occupied.

However, data from the FDAA shows that large-scale fires are relatively infrequent—only around 20–30 such incidents occur annually in Amsterdam. Moreover, a flexible dispatch model can still ensure that partial teams respond immediately for initial support, while full operations are coordinated as needed.

2.10 The Role of Simulation

Before implementing such a system, its theoretical benefits should be validated. Simulation plays a critical role in this process. We propose enhancing the fire department's existing dispatch simulator to model the new deployment strategies. By experimenting with various configurations and dispatch rules, we can evaluate their effects on overall response time, personnel utilization, and incident coverage.

In addition to scenario simulations, mathematical optimization models should also be incorporated into the simulator. Currently, the FDAA's operational model does not leverage algorithmic optimization in its dispatching process. However, recent studies within the FDAA have explored the use of optimization techniques for resource planning and decision-making. Integrating these models into the simulator will allow identification of optimal or near-optimal dispatch strategies dynamically during simulations, further improving the credibility and utility of the results. This enhancement can support the development of a semi-automated decision-support tool recommending optimal unit assignments under varying conditions.

If simulations demonstrate meaningful improvements over the current approach, the department will have a strong foundation to justify pilot testing and investment in more flexible deployment systems.

3

Literature Review

3.1 Simulation Tools in Emergency Service Planning

The use of simulation in emergency response planning has evolved into a critical method for evaluating and improving operational strategies without the risks and costs of real-world experimentation. The following sections review two foundational simulation systems that are relevant to this thesis: the TIFAR simulator developed for the Dutch ambulance sector, and the FDAA simulator, which represents the current state-of-the-art in fire dispatch simulation in the Netherlands. These prior works provide both methodological inspiration and a technical baseline for the simulation-integrated optimization framework.

3.1.1 Prior Simulation Research in the Netherlands

The development of dispatch simulation tools for emergency services in the Netherlands has a strong academic basis. Well before the FDAA's own simulation initiative, the TIFAR (Testing Interface For Ambulance Research) tool was created to evaluate ambulance dispatch strategies (van Buuren et al., 2012 (9)). TIFAR allowed Emergency Medical Service (EMS) managers to test different dispatch strategies and relocation policies under realistic constraints, using historical call and movement data.

The simulator was designed for the structure of the Dutch EMS, where each of the 25 regions operates its own ambulance provider (RAV), and incoming calls are classified by urgency: A1 (life-threatening), A2 (non-life-threatening), and B (pre-scheduled, non-urgent). Simulations with real demand and travel data showed that TIFAR's results closely matched actual performance for high-priority (A1) incidents, confirming the tool's accuracy. TIFAR also handled large geospatial datasets (over 32,000 vertices), supported real-time decision visualization, and allowed flexible rule and vehicle configurations.

Although TIFAR focused mainly on EMS, its success demonstrated the value of simulation in emergency logistics. This work helped establish simulation as a reliable method for testing what-if scenarios and optimizing emergency response systems in the Netherlands. The more recent FDAA simulator builds on this experience, focusing on fire response and allowing integration with mathematical optimization models.

3.1.2 The FDAA Simulator: Toward Intelligent and Data-Driven Fire Dispatch

A notable recent development in the FDAA's move toward a data-driven operational model is the simulation system created as part of a professional doctorate project titled *The Data-Driven Fire Department* (Van den Bogaert, 2019 (5)). This project aimed to provide the FDAA with decision-support tools to reduce emergency response times at strategic, tactical, and operational levels.

A key outcome was the creation of a verified and validated simulation software system. The simulator allows users to model and evaluate changes to station locations, vehicle and staff allocations, dispatch rules, and relocation strategies. Its modular design supports automated scenario comparisons, providing insight into how different resource configurations affect performance metrics such as response time and coverage. The study also introduced predictive models based on demographic data to improve incident forecasting and tested both rule-based and AI-based approaches for dynamic vehicle relocation. While deep reinforcement learning showed limited improvement over existing algorithms (MCRP+LBAP), the simulation-based response time distributions proved useful for evaluating relocation decisions.

The simulator developed in this project represents a foundational step toward more intelligent dispatch systems. However, it is not yet integrated with advanced optimization algorithms. Integrating mathematical models—especially optimization techniques—into the simulation framework could further enhance its decision-support capabilities. A more detailed discussion of the simulator's structure, functions, and planned improvements is provided in the Simulaton chapter of this thesis.

3.2 Evolution of Fire Resource Optimization Models

This section traces the development of fire resource allocation models—from early static location optimization to real-time redeployment heuristics and modern hybrid frameworks. Understanding this evolution provides critical context for the integration of simulation and optimization in today's fire dispatch systems.

3.2.1 Toregas et al. (1971): Static Fire Station Location Using Set Covering

Toregas et al. (10) introduced the Location Set Covering Model (LSCM), one of the earliest mathematical formulations for placing emergency service facilities. The model uses binary integer programming to determine the minimal number and placement of stations such that every demand point is within a pre-specified response radius. The model simplifies reality by assuming equal station costs, static demand, and full availability of units. Nevertheless, it laid the foundation for decades of facility location research in emergency services. It has since been extended to more flexible formulations, such as the Maximal Covering Location Problem (MCLP) (Church & ReVelle, 1974 (11)) and dynamic relocation models. The concepts from this foundational work still underpin many modern simulation-based and optimization-driven fire resource planning tools.

3.2.2 Kolesar and Walker (1974): Real-Time Unit Relocation During Major Incidents

While early models like Toregas et al. (10) focused on the static location of fire stations, real-world operations often require dynamic decision-making in response to unfolding events. A major advancement came from Kolesar and Walker (12), who addressed the issue of fire station coverage degradation during large-scale incidents. Their heuristic model supported real-time redeployment of available units to maintain city-wide response coverage when stations became empty. It was one of the first approaches to integrate computational logic into dispatch operations, balancing fairness, efficiency, and response speed under operational constraints. Their model showed that meaningful improvements in coverage during major incidents were possible, and the algorithm eventually became part of FDNY's real-time dispatch logic. It also established a precedent for integrating simulation and algorithmic decision-making into fire department operations—an idea central to the current project's vision.

3.2.3 MCRP + LBAP Model by FDAA

The FDAA proposed a novel relocation framework to address the operational demands of modern fire services, as described in Fire Truck Relocation During Major Incidents (Usanov et al., 2019 (3)). Their approach combines a Maximum Coverage Relocation Problem (MCRP) with a Linear Bottleneck Assignment Problem (LBAP), producing a heuristic relocation strategy that balances performance, fairness, and practical feasibility.

Compared to earlier models relying on complex cost functions or static redeployment rules, the FDAA framework:

- Uses historical incident distributions to inform relocation decisions;
- Introduces a tunable relocation "willingness" to control the number of vehicle movements based on operational capacity and dispatcher preferences;
- Maintains minimum coverage fairness across all city zones;
- Can be applied in real time with limited computational overhead.

During major incidents involving multiple units, the model identifies which idle trucks should be relocated and to which stations to restore balanced city-wide coverage. The optimization is solved dynamically, accounting for real-time unit availability and spatial risk patterns derived from historical data. The model was validated through extensive discrete-event simulations and calibrated using ten years of real incident and dispatch data from the FDAA. Compared to benchmark strategies—including static baselines and the Kolesar-Walker approach—the MCRP+LBAP heuristic significantly improves response reliability under high-demand scenarios. These features also make it suitable for frontline adoption by dispatch personnel.

Our project aims to embed and extend this model within a flexible simulation framework, enabling systematic evaluation of dispatch strategies under varied deployment configurations. The detailed design of the MCRP+LBAP algorithm, along with the adjustments made to adapt it to our simulator architecture, will be discussed in Chapter 5.

3.3 Other Research by FDAA

Beyond real-time dispatch and dynamic relocation, the FDAA has actively explored a range of optimization and forecasting strategies to enhance its operational planning. Two notable studies are summarized below to highlight the broader context of mathematical modeling within the organization.

3.3.1 Strategic Base Relocation and Crew Distribution Optimization

In collaboration with academic partners, the FDAA developed a mathematical programming model for optimizing fire station locations and vehicle allocations (Van den Berg et al., 2017 (2)). This model incorporates several real-world constraints:

- Some base stations must remain fixed due to historical or logistical reasons;
- Multiple vehicle types are supported, each with different response time targets;
- Differentiates between professional and volunteer firefighter crews.

Key findings from the study include:

- Relocating just 3 out of 19 stations reduced late arrivals by over 50%;
- It was unnecessary to add new stations—optimizing the locations of existing bases proved effective and more cost-efficient;
- Minor adjustments in crew or vehicle distribution raised coverage from 83.48% to 92.88%.

This model has already informed operational decisions, such as not replacing a specific base, and is currently in use in two other projects. It demonstrates that mathematical optimization models are not only theoretically valuable but also highly applicable in practice when carefully tailored to operational realities.

3.3.2 Weather-Informed Incident Forecasting Models

Beyond dispatch and station optimization, the FDAA has explored predictive modeling to improve resource planning under uncertain demand. One representative study focused on incident forecasting under severe weather conditions (Legemaate et al., 2021 (4)). The researchers developed a machine learning–based framework to forecast the number of incidents per fire station, incorporating not only historical incident data but also weather variables such as rain, wind, temperature, and visibility. Key findings include:

- Rain and wind strongly correlate with increased incident frequency, especially for storm-related calls;

- Temperature influences incidents in a non-linear fashion, while visibility was found to be the least impactful;
- Inhomogeneous Poisson processes outperformed simpler models for rare but severe incidents.

The best-performing model was an ensemble approach, which combined the strengths of multiple forecasting techniques. Importantly, the study also identified directions for future improvement, such as:

- Shifting from station-based to region-based forecasting, to better capture localized environmental risk;
- Incorporating seasonal vegetation factors into risk assessment;
- Clustering incidents by type for more targeted forecasts.

3.4 International Research on Fire Brigade Optimization

Beyond the FDAA, fire agencies worldwide have increasingly adopted mathematical optimization and data-driven modeling to enhance deployment efficiency. These international efforts illustrate the growing role of operations research and simulation as essential tools in modern emergency response planning. Two recent studies from Slovakia and Chile exemplify how different methodological approaches—ranging from integer programming to forecast-informed simulation—have been tailored to improve fire station placement, vehicle allocation, and response coverage.

3.4.1 Slovak Research on Fire Brigade Optimization

A study by Leitner et al. (2023 (13)) from the University of Žilina explores various mathematical programming approaches to optimize fire brigade deployment in Slovakia. The study presents a progression from basic to advanced models, aiming to balance efficiency and fairness in emergency coverage. The six models discussed include:

1. Basic Coverage Model (p-median) – Minimizes total distance between demand nodes and fire stations.
2. Radial Coverage Model – Focuses on binary coverage within a maximum distance threshold.

3. Generalized Disutility Model – Allows for more flexible distance penalties beyond simple thresholds.
4. Min–Max Fairness Model – Minimizes the worst-case distance to the nearest station, ensuring equity but possibly worsening overall averages.
5. Fair Coverage Model with Tolerance – Introduces parameters to control how many users can be outside a defined service radius.
6. Minimum Required Station Model – Calculates the minimal number of stations needed to satisfy fairness constraints.

The authors conclude that while simple models suffice for initial planning, advanced fairness-oriented models provide robustness and political feasibility. Full implementation, however, requires further simulation and real-world validation, particularly for assessing robustness and trade-offs between multiple objectives. This demonstrates the value of simulation tools for evaluating theoretical models in practical contexts.

3.4.2 Chilean Case Study: Fire Station and Fleet Optimization in Santiago

A example comes from the Santiago Fire Department (SFD) in Chile, where Pérez et al. (2016 (14)) developed an optimization model based on real operational data (over 22,000 emergency calls from 2007 to 2011, extended to nearly 38,000 after forecasting to 2015) to evaluate how fire station locations and fleet assignments impact emergency response performance. Several optimization experiments were conducted:

- Baseline scenario: Using existing station locations and vehicle assignments, establishing the reference performance level;
- Scratch (full redesign) scenario: Allowing unrestricted relocation of stations and full reallocation of vehicles, to estimate the upper bound of performance improvements;
- Fleet reallocation only: Keeping all station locations fixed, while optimizing vehicle deployment across stations;
- Single-station relocation: Allowing only one fire station to be relocated, with full fleet optimization. This represents a politically and operationally feasible middle ground.

The results were significant:

- Merely reallocating the fleet (no new stations) improved coverage by up to 10–35%;
- A single-station relocation combined with optimal fleet reallocation improved coverage by up to 30–50%;
- In terms of time-specific performance, the percentage of calls responded to within 6 minutes increased from 70% to 95% in some configurations—all without adding new vehicles or stations.

4

Data

The data used in this study were generated and processed according to the procedures described in the background section. Here, we provide a detailed overview of the original data provided by the FDAA and the additional information derived to support simulation-based analyses.

The FDAA provided five primary sources: Incident, Deployment, Station, Vehicle, and Deployment Characteristic. Except for Deployment Characteristic, these largely follow the structure used in *The Data-Driven Fire Department* (5), with updated records covering 2018 to 2025. This expanded data includes previously unavailable information from the GBA fire station, providing broader geographic and operational coverage.

Due to changes in data schema, naming conventions, and variable values from system updates, we performed preprocessing to remap the new format into a structure compatible with the simulator. These steps are detailed in Subsection 4.2.

On this updated data, we generated the additional inputs required by the simulator using the same methods as the original project, including Prophet-based forecasts, OSRM travel time matrices, and location datasets. Furthermore, the Deployment Characteristic data was analyzed to extract statistical distributions for incident escalation, which are central to modeling multi-stage emergencies.

4.1 Data Sets

The simulator in this study is based on historical incident data provided by the Amsterdam Fire Department. The core datasets include Incident, Deployment, Station, Vehicle, and

Deployment Characteristic, covering the period from May 2017 to April 2025. Together, these datasets combine previously available records with newly obtained data extending through 2025. Below, we describe each dataset in detail, outlining its structure, key attributes, and role in the simulation model.

Incident

The Incident dataset forms the foundation of all other records, containing information on incidents occurring between January 2008 and May 2025. For each incident, the fire department logs start and end times, type, priority level, and location. The dataset consists of 227,100 rows and 104 columns, with each row representing a single incident. This dataset merges records from the old and updated systems. The new system introduces additional variables and revised naming conventions, increasing the number of features from 56 to 104. To ensure compatibility with the simulator, we mapped the essential features back to the previous naming format without altering the original data content. These preprocessing steps are described in Subsection 4.2.

The simulator requires only 12 key features from this dataset, summarized in Table 4.1. Among these features, incident type is particularly important, as it determines the re-

Column Name	Column Type	Description
dim_incident_id	int64	Incident ID
dim_incident_incident_type	object	Type of incident
dim_datum_datum	object	Date
dim_prioriteit_prio	int64	Priority level
inc_dim_object_functie	object	Building type
dim_incident_start_datumtijd	object	Start time of incident
dim_incident_eind_datumtijd	object	End time of incident
dim_tijd_uur	object	Hour of the day
hub_vak_bk	float64	Hub area
kazerne_groep	object	Station group
st_x	object	X coordinate
st_y	object	Y coordinate

Table 4.1: Dataset schema

sources required and the temporal patterns of each event. The current classification system defines 18 types, including Outdoor fire, Alarm, Indoor fire, Resuscitation, Vehicle in water, and Aviation rescue, among others.

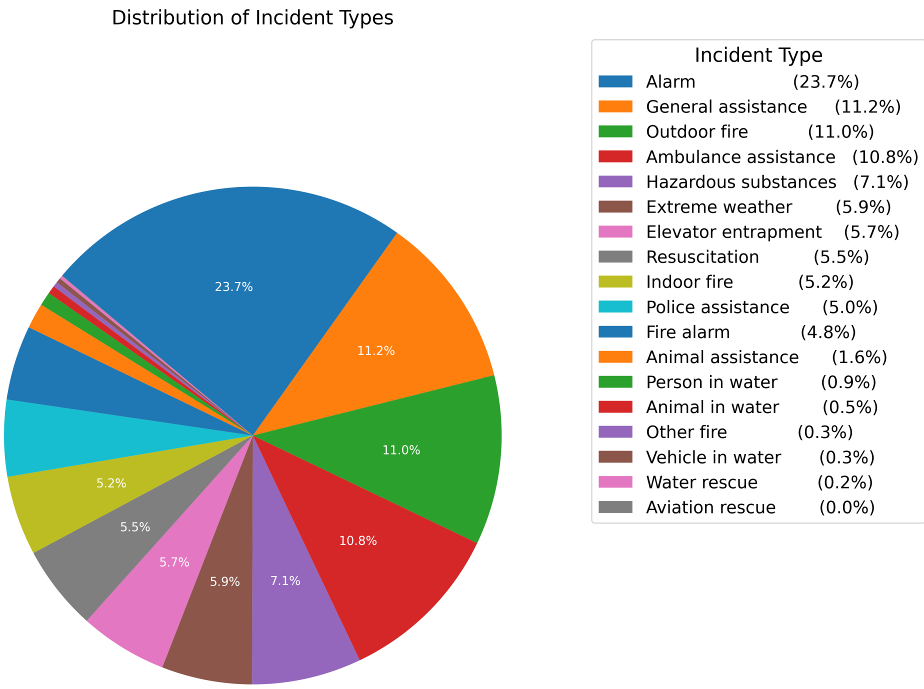

Figure 4.1: Pie chart showing proportions of incident types

As shown in Figure 4.1, Alarm incidents account for the largest share. These are typically triggered automatically and would often be categorized as Indoor fire if reported manually. Differences in reporting sources—automatic alarms versus human reports—also affect dispatch times. Human-reported incidents involve communication delays, while alarms often provide faster, though sometimes false, notifications. In our study, Alarm and Indoor fire incidents are modeled separately to capture these distinct dispatch characteristics.

All incident types are included in the simulation. It is important to note, however, that the Incident dataset records only one row per incident, without capturing changes in scale over time. To model the temporal evolution and escalation of incidents, we rely on the Deployment Characteristic dataset, which provides detailed records from May 2017 to April 2025. These data are used to fit distributions related to incident escalation, as discussed in later sections.

Deployment

The Deployment dataset records the dispatch details of vehicles for each incident, including vehicle type, departure and return times, and other operational timestamps. It contains 527,076 rows—each representing a single vehicle deployment—and 121 columns. Since multiple vehicles may be dispatched to the same incident, this dataset is substantially larger than the Incident dataset, with an average of 2.32 vehicles per incident. The data merges records from both the old and new systems, resulting in overlapping features with differing naming conventions. Necessary features have been mapped to ensure compatibility with the simulator (see Subsection 4.2).

Only 10 fields are required for the simulator, as listed in Table 4.2.

Column Name	Column Type	Description
hub_incident_id	int64	Associated incident ID
inzet_gealarmeerd_datumtijd	object	Time of alarm
inzet_uitgerukt_datumtijd	object	Time of departure
inzet_terplaatse_datumtijd	object	Time of arrival on scene
inzet_start_inzet_datumtijd	object	Start time of deployment
inzet_eind_inzet_datumtijd	object	End time of deployment
voertuig_groep	object	Vehicle group
kazerne_groep	object	Station group
inzet_terplaatse_volgnummer	int64	Arrival sequence number
inzet_rijtijd	object	Travel time

Table 4.2: Fields from the Deployment dataset required by the simulator.

The variable hub_incident_id links each deployment to its corresponding incident (dim_incident_id) in the Incident dataset. Based on these variables, the following operational metrics are derived:

- `Dispatch time` = Alarm time − Incident start time
- `Turnout time` = Departure time − Alarm time
- `Travel time` = Travel time
- `Response time` = Dispatch time + Turnout time + Travel time

Station and Vehicles

The Station dataset is relatively small, containing geospatial information about fire stations, including station name, address, postal code, city, longitude, and latitude. The Vehicle dataset provides the staffing and equipment structure of each station, recording the number of vehicles of each type and associated crew sizes. For confidentiality reasons, station names are represented by letters (see Table 4.3).

Station	TS	RV	HV	WO	TS_crew_ft	TS_crew_pt	RVHV_crew_ft	RVHV_crew_pt	WO_crew_ft	WO_crew_pt
A	2	1	1	1	1	1	1	0	0	1
B	1	1	0	1	1	0	1	0	1	0
C	1	0	0	0	1	0	1	0	0	0
D	1	1	0	0	1	0	1	0	0	0
E	1	1	0	0	1	0	1	0	0	0
F	1	0	0	0	1	0	1	0	0	0
G	2	0	0	0	0	2	0	0	0	0
H	1	1	0	0	1	0	1	0	0	0
I	1	1	0	0	1	0	1	0	0	0
J	1	0	0	0	0	1	0	0	0	0
K	1	0	0	0	1	0	1	0	0	0
L	1	0	0	0	1	0	0	0	0	0
M	1	1	1	0	1	0	2	0	0	0
N	1	1	0	0	1	0	1	0	0	0
O	2	0	0	0	0	2	0	0	0	0
P	2	0	1	0	1	1	1	0	0	0
Q	1	0	0	0	0	1	0	0	0	0
R	1	0	0	0	1	0	0	0	0	0

Table 4.3: Vehicle availability and staffing per station.

The dataset records the number of *crews*, not individual firefighters. The crew sizes are standardized as follows:

- TS: 6 firefighters
- RV and HV: 2 firefighters (shared crew)
- WO: 4 firefighters

Firefighter types are further distinguished as full-time (ft) and part-time/volunteer (pt). Subsequently, crew units are converted into individual headcounts, and the dataset is expanded to include additional vehicle types and crew compositions for the designed simulation scenarios.

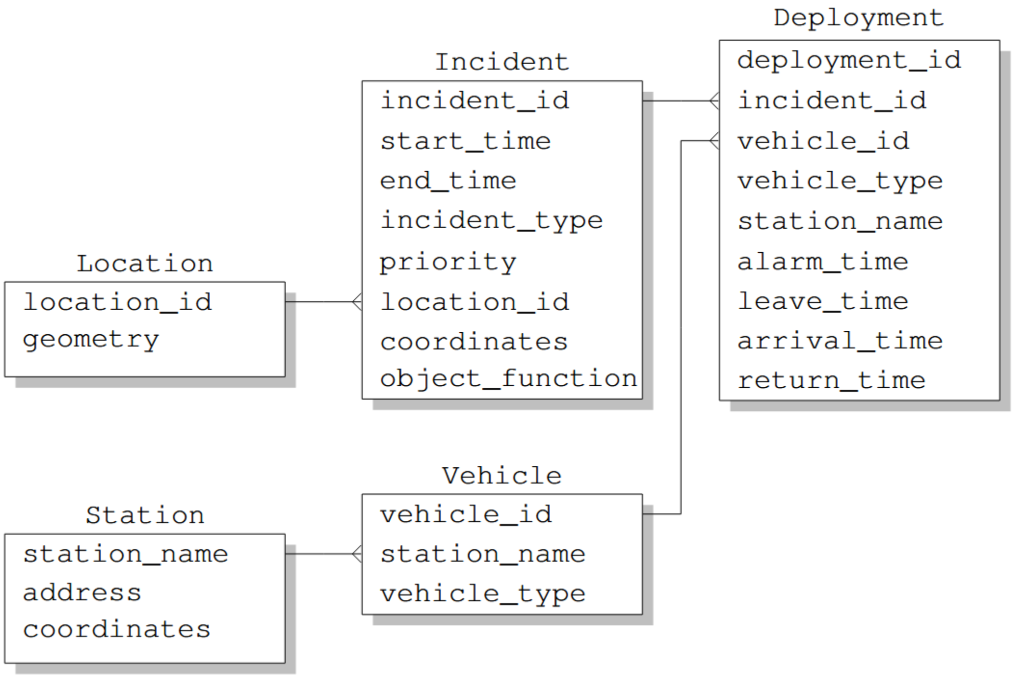

Figure 4.2: : Entity-relationship diagram (adapted from (5))

The relationships between these datasets are consistent with those presented in *The Data Driven Fire Department* (5). Simulation locations are generated using hub_vak_bk from the Incident dataset, enabling the creation of a synthetic Location dataset based on real incident locations.

Deployment Characteristic

A key enhancement in our simulator, compared to the simulator of *The Data Driven Fire Department* (5), is the inclusion of incident escalation dynamics. The Deployment Characteristic dataset captures how the scale of an incident evolves as additional resources are dispatched. It contains 61,615 rows and 11 features, with the three most relevant fields being:

- `INCIDENT_ID`: Link to the Incident dataset
- `ACTUELE_KAR_WAARDE`: Final registered scale of the incident
- `DTG_REG_WAARDE`: Datetime of final scale registration

Incident scales are categorized into four levels: small, medium, major, and very major. The dataset records only the final scale for each incident. We assume all incidents begin as small and escalate sequentially, although intermediate escalation points (e.g., small to medium) are not documented. For such cases, expert judgment is used to define plausible escalation timelines, which will be described in detail in Chapter 6.

While the inclusion of escalation processes enhances the realism of simulated incidents, the limited temporal coverage of this dataset (2017–2025) constrains precise statistical fitting for all scenarios. The primary aim is to demonstrate the feasibility and structure of this simulator enhancement, rather than to exactly replicate historical data. As the fire department continues to collect more detailed records, future work can leverage richer datasets for more accurate calibration.

4.2 Data Processing: Mapping and Data Generation

This section describes the preprocessing steps applied to harmonize and enrich the datasets used in the simulation. The procedures consist of two main components: (1) mapping existing features and values to ensure compatibility with *The Data Driven Fire Department* dataset (5), and (2) generating new datasets to support simulation and forecasting.

Mapping

Although the newly acquired data (2018–2025) shares structural similarities with the original dataset, there are notable differences in feature names and value formats. To ensure compatibility with the simulator, we mapped the new dataset fields to match the original format.

We adopted the variable naming conventions of the old dataset and mapped the corresponding new variables. The mappings for the Incident dataset are listed in Table 4.4. Similarly, the new Deployment dataset fields were mapped to the old format as shown in Table 4.5.

In addition to field-level mapping, we standardized incident types to align with the simulator-compatible categories, as summarized in Table 4.6.

This standardized mapping ensures consistent recognition of incident types by the simulator.

New Field Name	Mapped to (Old Format)
NR_Incident	dim_incident_id
Incident_Classificatie	dim_incident_incident_type
dim_incident_start_datumtijd	dim_datum_datum
Incident_Prio	dim_prioriteit_prio
(Missing in new data)	inc_dim_object_functie (handled in Simulator section)
INC_START_BRW_YYMMDD_HHMMSS	dim_incident_start_datumtijd
INC_EIND_BRW_YYMMDD_HHMMSS	dim_incident_eind_datumtijd
(Derived from end time)	dim_tijd_uur
Vak	hub_vak_bk
Kazerne_Groep	kazerne_groep
(Derived from “Vak” coordinates)	st_x, st_y

Table 4.4: Mapping of new Incident dataset fields to the old format.

New Field Name	Mapped to (Old Format)
NR_INCIDENT	hub_incident_id
Alarm_Tijd_TS	inzet_gealarmeerd_datumtijd
Uitruk_Tijd_TS	inzet_uitgerukt_datumtijd
Ter_Plaatse_Tijd_TS	inzet_terplaatse_datumtijd
Opdracht_Inzet_Tijd_TS	inzet_start_inzet_datumtijd
INZ_INCIDENT_EIND_TIJD_TS_HHMMSS	inzet_eind_inzet_datumtijd
Voertuig_Soort_Code	voertuig_groep
Kazerne_Groep	kazerne_groep
Voertuig_Rank	inzet_terplaatse_volgnummer
Rijtijd	inzet_rijtijd

Table 4.5: Mapping of new Deployment dataset fields to the old format.

Generated Data

To support simulation and forecasting, several datasets were generated following the methodology described in *The Data Driven Fire Department* (5):

- `prophet_forecast` — Generated using Facebook's Prophet model (15) to estimate the hourly arrival rate of incidents. The arrival rate varies by incident type, which is taken into account during model training.

- `OSRM_time_matrix` — Computed using the Open Source Routing Machine (OSRM) to estimate travel times between all pairs of relevant locations. As FDAA's coverage

Old Category	**Mapped to (New Category)**
Letsel eigen personeel	Hulpverlening algemeen (General assistance)
Hulpverlening algemeen	Hulpverlening algemeen
Beknelling / bevrijding	Hulpverlening algemeen
Buitensluiting	Hulpverlening algemeen
Brandbare vloeistoffen	Gevaarlijke stoffen (Hazardous substances)
Brandbare gassen	Gevaarlijke stoffen
Overige gevaarlijke stoffen	Gevaarlijke stoffen
Meten / overlast / verontreiniging	Gevaarlijke stoffen
Assistentie Ambulance	Assistentie ambulance (Ambulance assistance)
Afhijsen spoed	Assistentie ambulance
Hulpverlening algemeen Dieren	Hulpverlening dieren (Animal assistance)
Hulpverlening Dieren	Hulpverlening dieren
Binnenbrand	Binnenbrand (Indoor fire)
Buitenbrand	Buitenbrand (Outdoor fire)
OMS / automatische melding	Alarm
Brandgerucht / nacontrole	Brandgerucht (Fire alarm)
Liftopsluiting	Liftopsluiting (Elevator entrapment)
Assistentie Politie	Assistentie politie (Police assistance)
Persoon te water	Persoon te water (Person in water)
Voertuig te water	Voertuig te water (Vehicle in water)
Hulpverlening water algemeen	Hulpverlening water (Water rescue)
Storm en Waterschade	Extreem weer (Extreme weather)
Dier te water	Dier te water (Animal in water)
Reanimeren	Reanimatie (Resuscitation)

Table 4.6: Mapping of old incident type categories to standardized simulator-compatible categories.

area has expanded, new demand points were added, and the travel-time matrix was recalculated to reflect the updated geography.

Creating the Location Dataset

Location information, recorded as numeric Hub codes in the deployment dataset, was extracted and used to generate travel times via OSRM. The resulting Location dataset includes each point's longitude, latitude, and corresponding `st_x` and `st_y` coordinates for visualization.

This dataset is fundamental for defining:

- The service area of each fire station (based on OSRM travel times), and
- The response neighborhood used for candidate relocation zones in simulations.

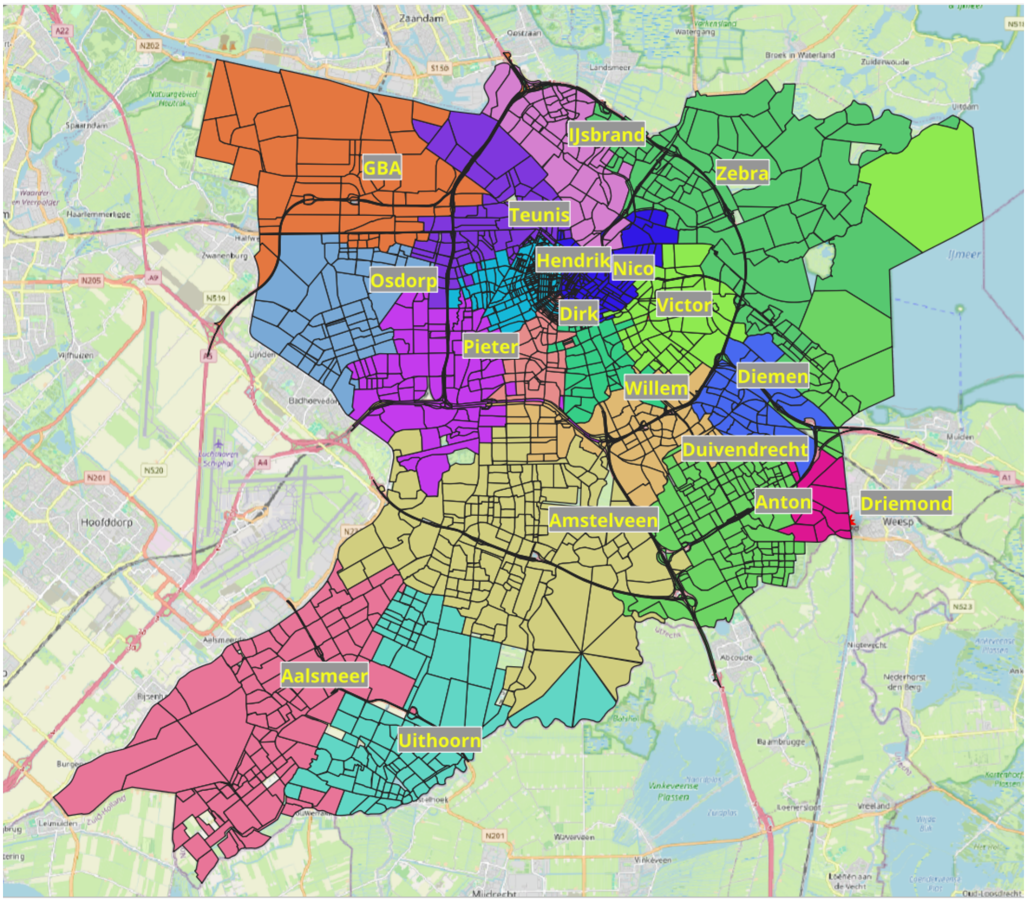

Figure 4.3: Figure: Service area maps generated based on OSRM travel times

As shown in Figure 4.3, the defined service areas correspond specifically to TS-series vehicles (TS), which are present at every station. Specialized vehicles, such as HV, RV, or WO, are not deployed uniformly, resulting in different service areas for different vehicle types. In this study, we focus on TS vehicles, which form the core operational units across all stations.

While these OSRM-based service areas provide a practical approximation, they may not fully reflect actual operations. For example, some nearby stations may not be the primary responders due to terrain constraints or prior experience, and certain stations have historically served specific regions regardless of distance. Nonetheless, for simulation purposes, these service areas are sufficiently accurate.

Finally, algorithmically generating service areas enhances user autonomy, allowing future simulator users to update service boundaries without manual requests from the fire department, thereby simplifying experimentation and deployment.

4.3 Data Analyse

Response time

The recording of response times is partially automated and partially reliant on manual input by firefighters. When the control center receives an emergency report, the system automatically logs the incident start time. Likewise, the moment the alarm is triggered at the fire station is recorded automatically. However, subsequent timestamps—such as vehicle departure, arrival at the scene, and return—must be manually logged by firefighters using a button inside the vehicle. Delays or omissions in manual input can introduce discrepancies between recorded and actual times.

For our analysis and distribution fitting, we adopt the filtering approach used in *The Data Driven Fire Department* (5), excluding dispatch or turnout times exceeding 10 minutes.

Dispatch time Dispatch time is entirely determined by the control center and is therefore independent of the incident location or the proximity of fire stations. Nonetheless, different incident types may exhibit distinct dispatch time distributions.

We focus on three categories: all incidents, Indoor Fire incidents, and Resuscitation incidents. Domain experts report that dispatch times typically center around 90 seconds.

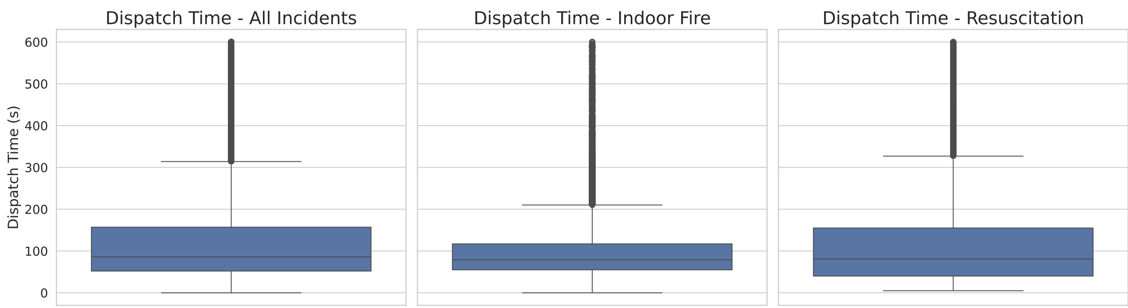

Figure 4.4: Boxplot comparison of dispatch times for all incidents, Indoor Fire, and Resuscitation.

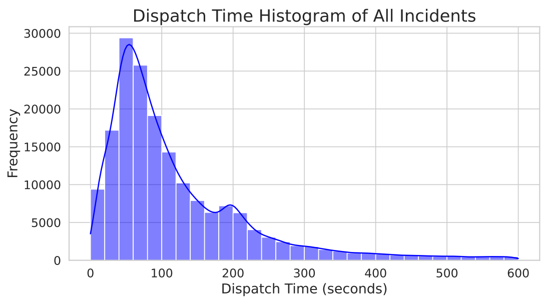

(a) All Incidents

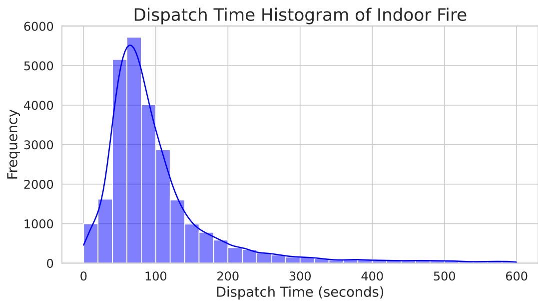

(b) Indoor Fire

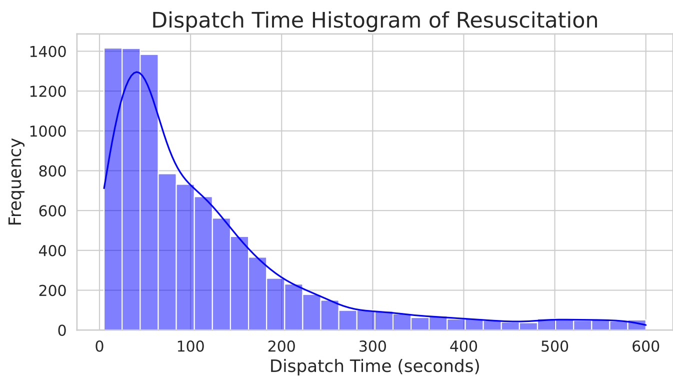

(c) Resuscitation

Figure 4.5: Dispatch time histograms for three incident types.

The boxplots in Figure 4.4 confirm that the median dispatch times for all three categories hover around 90 seconds. However, the histograms in Figure 4.5 reveal distinct distributional patterns, suggesting that dispatch times should be fitted separately for each incident type.

We adopt the Bayesian approach described in the article of Oliphant (2006 (16)) to calculate 95% confidence intervals for the mean and standard deviation of each category. Incident types with sufficient sample sizes are fitted individually; otherwise, the overall distribution across all incidents is used. This methodology is consistently applied to all time-related distributions in our study. The filtered dispatch time data ($\leq$ 10 minutes) is fitted with a log-normal distribution, as shown in Figure 4.6.

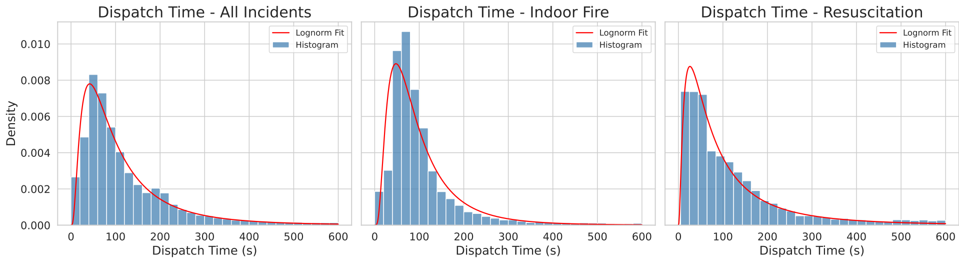

Figure 4.6: Histogram of dispatch time with fitted log-normal distribution.

Turnout time Turnout time is primarily influenced by firefighter type, incident priority, and vehicle type. Here, we analyze turnout times for pumper vehicles responding to priority-1 Indoor Fire incidents.

Most stations are staffed by full-time firefighters. Exceptions include Driemond and Duivendrecht stations, which are fully volunteer-staffed. Three additional stations—Aalsmeer, Uithoorn, and Diemen—operate with full-time staff on weekdays but switch to volunteer staffing on weekends. For analysis purposes, weekend data from these hybrid stations is treated as volunteer-staffed. Turnout times shorter than 30 seconds or exceeding 10 minutes are excluded to account for errors, such as accidental button presses before vehicle departure. As expected, volunteer-staffed stations exhibit longer turnout times compared to full-time stations (Figures 4.7–4.8). Figure 4.9 further illustrates the differences across individual stations.

Distributions are fitted based on station type, incident priority, and vehicle type. Non-priority-1 incidents are modeled jointly as non-urgent. The primary goal is to realistically capture differences in turnout behaviour between full-time and volunteer firefighters. If at least one volunteer firefighter is dispatched, the distribution fitted to volunteer-staffed responses is used.

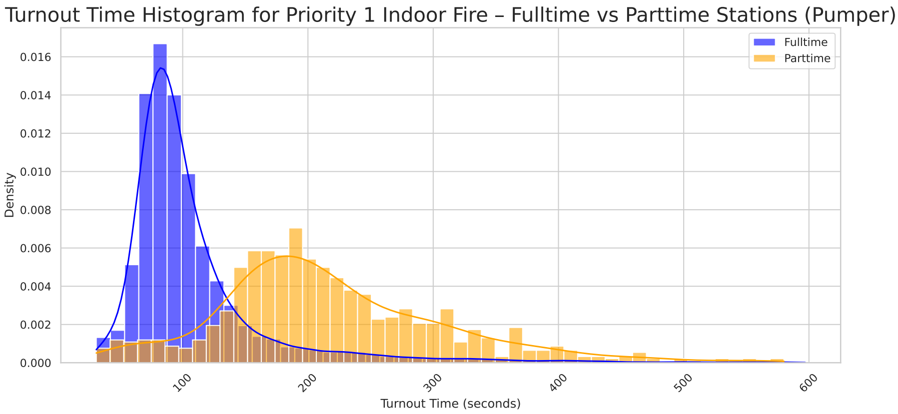

Figure 4.7: Turnout time histogram for priority 1 Indoor Fire – fulltime vs parttime stations (pumper).

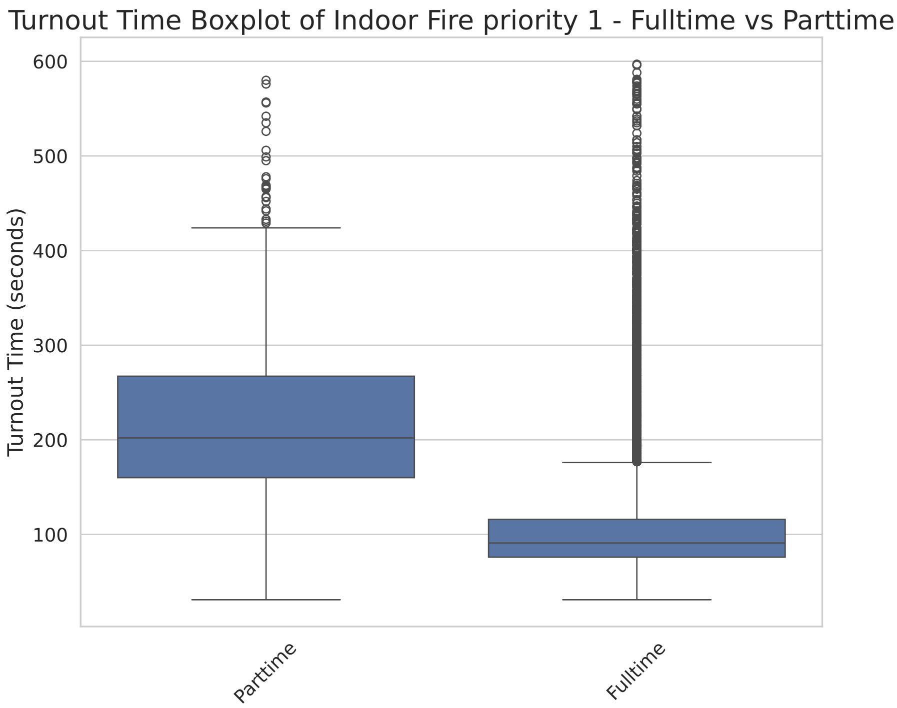

Figure 4.8: Turnout time boxplot of Indoor Fire priority 1 - fulltime vs parttime.

Figure 4.10 shows gamma distributions fitted for full-time and volunteer firefighters responding to priority-1 Indoor Fire incidents (pumper vehicles only).

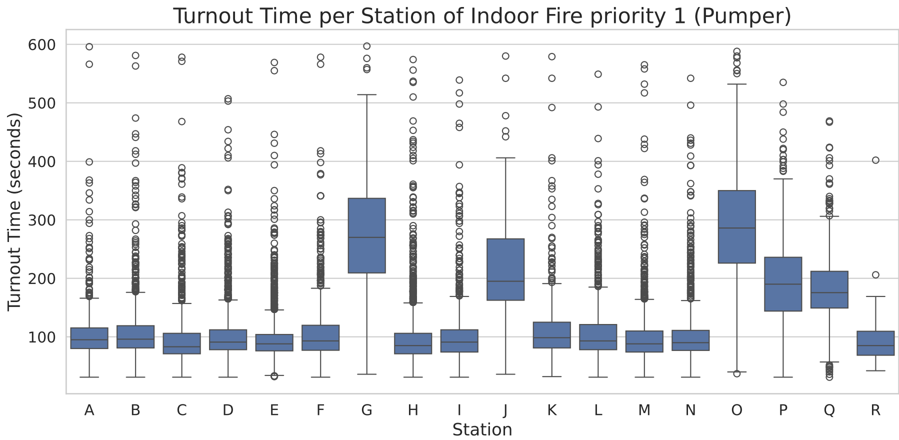

Figure 4.9: Turnout time per station of Indoor Fire priority 1 (pumper).

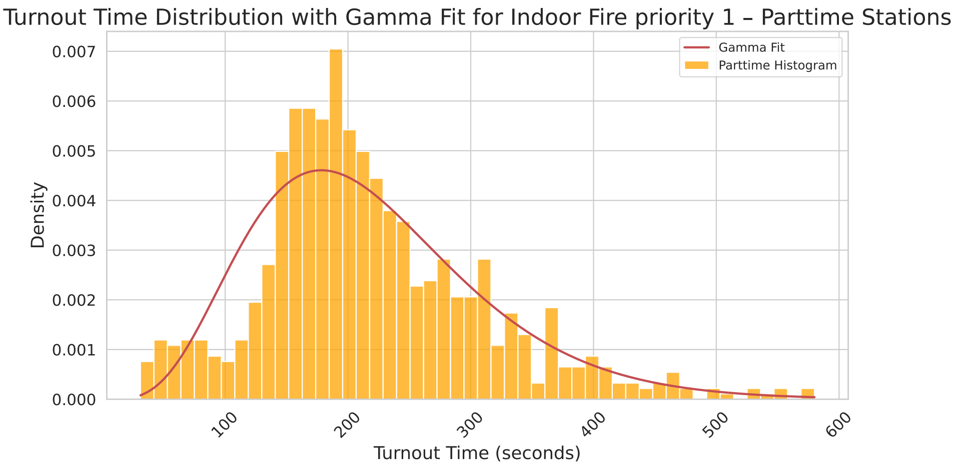

(a) Part-time

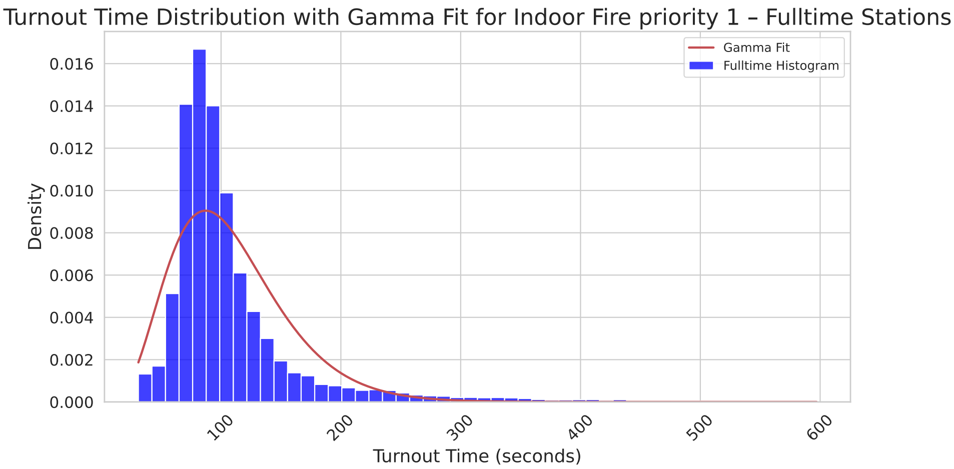

(b) Fulltime

Figure 4.10: Gamma distributions fitted for turnout times of volunteer and full-time firefighters responding to priority-1 Indoor Fire incidents (pumper vehicles only).

Travel time For travel times, we follow the approach used in *The Data Driven Fire Department* (5). OSRM-estimated travel times between locations are used as a baseline, while residuals from actual data capture variability. We use the following formulas:

$$\begin{aligned}
(a, b) &\sim \text{LinearRegression}\big(\text{TRAVEL}_{\text{est}}, \text{TRAVEL}_{\text{true}}\big), \\
\varepsilon_{\text{factor}} &= \frac{\text{TRAVEL}_{\text{true}} - a}{b \times \text{TRAVEL}_{\text{est}}}, \\
\varepsilon &\sim \text{Lognormal}\big(\mu, \sigma\big) \quad \text{fitted from } \varepsilon_{\text{factor}}, \\
\text{TRAVEL} &= a + b \times \varepsilon \times \text{TRAVEL}_{\text{est}}.
\end{aligned}$$

This method is applied separately for each vehicle type. It is important to note that OSRM travel estimates are based on regular vehicle speeds, while emergency vehicles generally operate at higher speeds. Moreover, the upcoming simulation framework in our study introduces new vehicle types such as MOTO, TS2, and TS4, which require further consideration. The treatment of these vehicles will be discussed in detail in Chapter 6.

Upgrading probability and incident time distribution

In the original simulator, developed by *The Data Driven Fire Department* (5), incident duration was modeled using historical on-scene times of dispatched vehicles. However, the scale or severity of an incident was not explicitly represented. Instead, the simulator implicitly captured scale by determining the required number of vehicles from historical data. For example, an Indoor Fire requiring one truck was treated as small, whereas one requiring three or more trucks was considered medium or large. This approach implied that large-scale fires occurred immediately. In reality, however, most fires begin small and may escalate over time.

In contrast, we propose a design in which incidents always start as small and may escalate step by step. This better reflects real-world dynamics, where fires typically start small and escalate only if uncontrolled. To support this mechanism, we analyze historical data to derive upgrading probabilities and fit duration distributions for each severity level. Each level is assigned its own time distribution. In this setup, the original simulator's on-scene time is replaced by fitted distributions corresponding to the current level of the incident. This requires the Deployment Characteristic dataset.

Upgrading Probability In our model, incidents always escalate sequentially from one level to the next. Although we only observe the final scale of an incident, intermediate upgrading probabilities can be derived using conditional frequency ratios:

$$\begin{aligned}
p_{\text{small}\to\text{medium}} &= \frac{f_{\text{medium}} + f_{\text{major}} + f_{\text{very major}}}{f_{\text{total}}}, \\
p_{\text{medium}\to\text{major}} &= \frac{f_{\text{major}} + f_{\text{very major}}}{f_{\text{medium}} + f_{\text{major}} + f_{\text{very major}}}, \\
p_{\text{major}\to\text{very major}} &= \frac{f_{\text{very major}}}{f_{\text{major}} + f_{\text{very major}}}.
\end{aligned}$$

For the **Indoor Fire** category, we observed:

- $f_{\text{small}} = 4722$
- $f_{\text{medium}} = 2113$
- $f_{\text{major}} = 505$
- $f_{\text{very major}} = 410$

This results in:

- $p_{\text{small}\to\text{medium}} = \dfrac{2113 + 505 + 410}{7750} \approx 0.39$
- $p_{\text{medium}\to\text{major}} = \dfrac{505 + 410}{2113 + 505 + 410} \approx 0.30$
- $p_{\text{major}\to\text{very major}} = \dfrac{410}{505 + 410} \approx 0.45$

The Deployment Characteristic dataset provides the final scale (`ACTUELE_KAR_WAARDE`) but only for recent years. Consequently, this dataset is smaller than the main Incident dataset, limiting the analysis for now, though the fitting quality will improve as more data accumulates. Probabilities are fitted per incident type; if sample sizes are insufficient, we fall back on aggregate distributions across all incidents. Not all incident types have multiple levels of severity—for example, Resuscitation and Elevator Entrapment incidents do not escalate, and their durations continue to be modeled using fitted on-scene times.

Incident time distribution We next analyze the timing of upgrades and total duration for each severity level. Incident IDs are linked from the Incident dataset to obtain start and end times, while the final severity registration time (`DTG_REG_WAARDE`) is used to calculate duration metrics. For clarity, we define:

- **Start moment**: when the incident was first reported
- **Upgrade moment**: when a new severity level was registered
- **End moment**: when the incident concluded

Two types of duration distributions are fitted:

- **Upgrade Time**: **Upgrade moment – Start moment**
- **Total Duration**: **End moment – Start moment**

The first is used to simulate the timing of upgrades (and subsequent dispatch of additional vehicles), while the second defines the full duration of an incident, after which vehicles return.

In theory, these distributions may differ by incident priority. However, due to data limitations, we do not segment by priority and instead use all available data for each incident type. We filter out durations shorter than 1 minute or longer than 24 hours. For upgrade times specifically, we additionally remove values above 10 hours.

Figures 4.11–4.14 illustrate log-normal fits for all incidents and indoor fire incidents, segmented by severity level.

From the upgrade time distributions, we observe that some incidents are registered as large-scale immediately—for instance, chemical factory fires may escalate rapidly, requiring dispatch of multiple vehicles across all levels within a short time.

In the total duration distributions, some high-scale incidents exhibit unexpectedly short durations. While this may reflect efficient containment, it could also arise from cases initially declared resolved but later re-ignited. Although we filter out durations shorter than 1 minute, the underlying causes of such anomalies cannot be determined with certainty given the data limitations.

Some incident types and severity levels lack sufficient data to support reliable distribution fitting. However, we have made use of all available data to the best extent possible. As more data becomes available in the future, these fits can be refined further.

Since only the final severity level is recorded, the actual timing of intermediate upgrades remains unknown. To address this, we introduce upgrade ratios based on domain expertise. If an incident is ultimately labeled as “major” or “very major,” we simulate intermediate

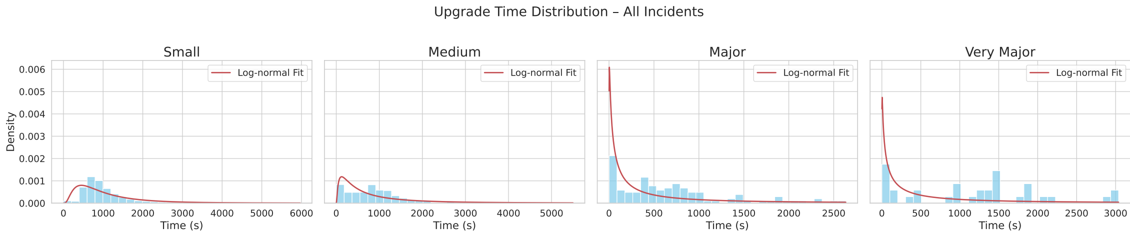

Figure 4.11: Upgrade time distribution – All Incidents

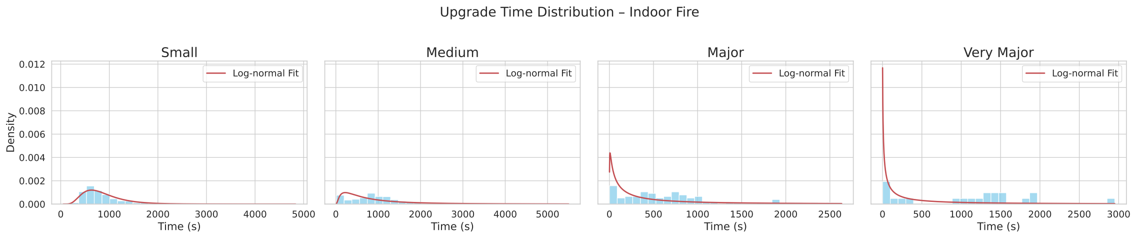

Figure 4.12: Upgrade time distribution – Indoor Fire

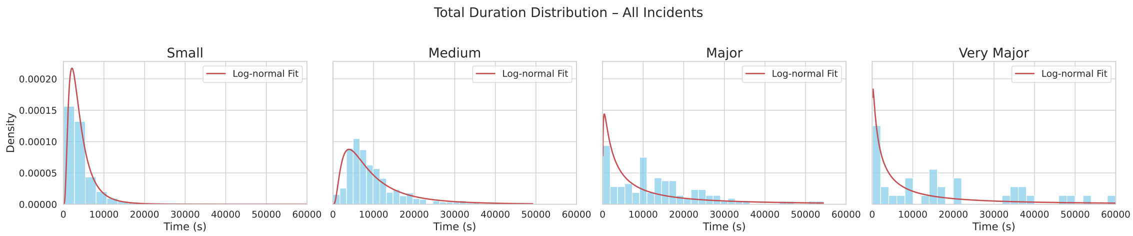

Figure 4.13: Total Duration distribution – All Incidents

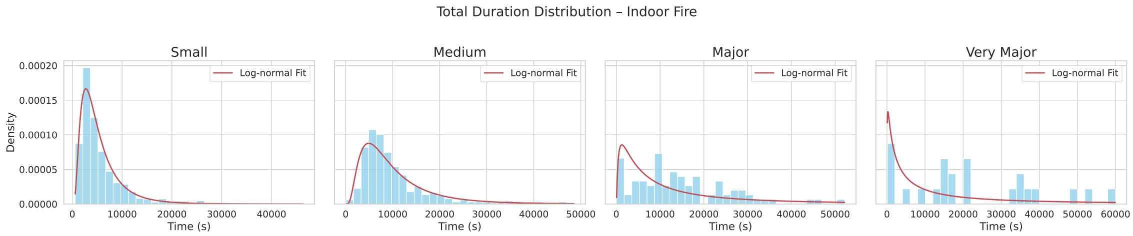

Figure 4.14: Total Duration distribution – Indoor Fire

upgrades by distributing the total upgrade time across levels using these ratios. The details of this mechanism are explained in Chapter 6.

5

Algorithm

This chapter describes the two algorithms used in the study. The first is the relocation algorithm, a combination of the MCRP and the LBAP (3), applied to move vehicles between stations when required. The second is the Station Dispatch Optimization Model (SDOM), developed in this work to optimize vehicle and crew deployment within a station.

Our experiments focus on applying the relocation algorithm to TS6 and Indoor Fire incidents, as these have the greatest operational impact. The approach, however, can be extended to other vehicle types and incident categories. The SDOM is mainly designed for dispatching TS-series vehicles (MOTO, TS2, TS4, and TS6).

5.1 Relocation Algorithm

During major incidents, multiple trucks are dispatched and may remain on-site for extended periods. As a result, the originating stations can run low on available vehicles, reducing their coverage. If new incidents occur nearby, vehicles must be sent from more distant stations, increasing response times and the risk of injuries and property damage.

To mitigate this, idle vehicles can be temporarily relocated from other stations to maintain regional coverage. The MCRP+LBAP algorithm determines how to relocate fire trucks during major incidents to maximize overall coverage while minimizing travel time.

The approach consists of two parts:

- **MCRP** — Determines which vehicles should be relocated from which stations to which destination stations, with the objective of maximizing coverage. The output is a set of origin-destination pairs indicating the required relocations.

- **LBAP** — Refines the MCRP solution by assigning specific relocation pairs between origins and destinations to minimize total travel time. While the MCRP maximizes coverage without considering travel time, LBAP ensures the relocation plan achieves minimal travel time.

For example, if the MCRP suggests sending one vehicle from station A to B and another from C to D, but A is closer to D and C is closer to B, LBAP will swap assignments to reduce travel time.

MCRP Model Formulation

Parameters and Sets:

- $\mathcal{N}$: set of stations
- $\mathcal{K}_n = \{\mathrm{RN}(N) \mid |N| = n\}$: set of response neighbourhoods, where each neighbourhood consists of the n closest stations to a given demand location
- A_n: incidence matrix with $a_{ik}^n = 1$ if station $i \in \mathcal{N}$ serves neighbourhood $k \in \mathcal{K}_n$, and 0 otherwise
- f_i: number of trucks available at station i immediately after a major incident
- d_i: total demand in the service area of station i
- $W \in [0, 1]$: weight parameter balancing the two objectives (coverage benefit vs. relocation cost)
- $\mathcal{E} = \{i \in \mathcal{N} \mid f_i = 0\}$: set of empty stations
- $S = \{i \in \mathcal{N} \mid f_i = 1\}$: set of stations with exactly one available truck
- $\mathcal{M} = \{i \in \mathcal{N} \mid f_i \geq 2\}$: set of stations with more than one truck

Decision Variables:

- $x_{ij} \in \{0, 1\}$: equals 1 if a truck is relocated from station i to station j, 0 otherwise
- $z_i \in \{0, 1\}$: equals 1 if station i becomes empty after relocation
- $y_i \in \mathbb{Z}_{\geq 0}$: number of trucks at station i after relocation

Objective Function:

$$\begin{aligned} \max \quad & W \left(\sum_{i \in S} \sum_{j \in \mathcal{E}} x_{ij}(d_j - d_i) + \sum_{i \in \mathcal{M}} \sum_{j \in \mathcal{E}} x_{ij} d_j - \sum_{i \in \mathcal{M}} z_i d_i \right) \\ & - (1 - W) \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{E}} x_{ij} \end{aligned} \tag{1}$$

Constraints:

$$\sum_{i \in \mathcal{N}} a_{ik}^{n} y_i \geq 1, \qquad \forall k \in \mathcal{K}_n \qquad \text{(coverage)} \tag{2}$$

$$\sum_{j \in \mathcal{N}} x_{ij} \leq f_i, \qquad \forall i \in \mathcal{N} \qquad \text{(supply limit)} \tag{3}$$

$$\sum_{j \in \mathcal{N}} x_{ji} \leq 1, \qquad \forall i \in \mathcal{E} \qquad \text{(one truck per empty)} \tag{4}$$

$$1 - z_i \leq y_i, \qquad \forall i \in \mathcal{M} \qquad \text{(emptiness logic)} \tag{5}$$

$$y_i = f_i + \sum_{j \in \mathcal{N}} x_{ji} - \sum_{j \in \mathcal{N}} x_{ij}, \qquad \forall i \in \mathcal{N} \qquad \text{(truck count update)} \tag{6}$$

$$x_{ij} = 0, \qquad \forall i \in \mathcal{N},\ \forall j \in S \cup \mathcal{M} \qquad \text{(relocation only to empty)} \tag{7}$$

$$x_{ij}, z_i \in \{0, 1\}, \qquad \forall i, j \in \mathcal{N} \qquad \text{(binary variables)} \tag{8}$$

$$y_i \in \mathbb{Z}_{\geq 0}, \qquad \forall i \in \mathcal{N} \qquad \text{(non-negative integer)} \tag{9}$$

Constraint Descriptions:

- (2) Each response neighbourhood must be covered by at least one truck after relocations
- (3) No station can send more trucks than it originally has
- (4) Each empty station can receive at most one truck
- (5) If a multi-truck station becomes empty, z_i must be 1
- (6) Tracks truck counts before and after relocations
- (7) Relocations can only go to empty stations
- (8) Binary indicators for relocation and emptiness
- (9) Truck counts must be non-negative integers

LBAP Model Formulation

Parameters:

- O: Set of origin stations (stations sending out trucks)
- D: Set of destination stations (stations receiving trucks)
- $t_{g(o)g(d)}$: Travel time from the location of origin station $o \in O$ to the location of destination station $d \in D$, where $g(\cdot)$ maps a station to its corresponding demand location

Decision Variables:

- $\widehat{x}_{od} \in \{0,1\}$: Equals 1 if a truck is relocated from station o to station d, and 0 otherwise

Objective Function:

$$\min \quad \max_{o \in O,\, d \in D} t_{g(o)g(d)}\, \widehat{x}_{od} \tag{10}$$

Constraints:

$$\sum_{d \in D} \widehat{x}_{od} = 1, \qquad \forall o \in O \qquad \text{(each origin assigns one truck)} \tag{11}$$

$$\sum_{o \in O} \widehat{x}_{od} = 1, \qquad \forall d \in D \qquad \text{(each destination receives one truck)} \tag{12}$$

$$\widehat{x}_{od} \in \{0,1\}, \qquad \forall o \in O, \forall d \in D \qquad \text{(binary variable)} \tag{13}$$

Constraint Interpretations:

- (11) Each origin station must relocate exactly one truck
- (12) Each destination station must receive exactly one truck
- (13) Assignments are binary

5.2 Personnel Considerations in Relocation Algorithm

In the original relocation algorithm, f_i represents the number of available vehicles at station i. Although personnel were not explicitly modeled, it was assumed that each available vehicle had a full crew (6 firefighters). A vehicle was considered available only if the station had a complete team ready.

In our revised scenario, vehicles do not need to be fully staffed at the origin. Instead, the combined personnel at the origin and destination must suffice to form a complete team. The relocation algorithm is thus adjusted to account for personnel allocation.

For example, if destination station B has 6 available firefighters but no vehicles, sending just one firefighter from station A to drive a relocated vehicle completes the team at B. This keeps more personnel at A available for other incidents.

Additional Parameters and Decision Variables

- c_i: number of firefighters available at station i
- t: team size required for a complete team (6 in this study)
- $\tilde{x}_{ij}$: number of firefighters dispatched from station i to station j

New Constraints:

$$
\begin{aligned}
x_{ij} \cdot t &\leq \tilde{x}_{ij} + c_j, && \forall i \in \mathcal{N},\ j \in \mathcal{N} && (14)\\
\tilde{x}_{ij} &\leq x_{ij} \cdot c_i, && \forall i \in \mathcal{N},\ j \in \mathcal{N} && (15)\\
\tilde{x}_{ij} &\geq x_{ij} \cdot \max(1, t - c_j), && \forall i \in \mathcal{N},\ j \in \mathcal{N} && (16)\\
\sum_{j \in \mathcal{N}} \tilde{x}_{ij} &\leq c_i, && \forall i \in \mathcal{N} && (17)
\end{aligned}
$$

Constraint Interpretations

- **(14)**: Relocation from station i to station j is allowed only if the dispatched firefighters from station i, combined with those available at station j, are sufficient to form a complete team of size t.
- **(15)**: The number of firefighters dispatched from station i cannot exceed the available personnel at that station.
- **(16)**: At least $\max(1, t - c_j)$ firefighters must be dispatched from station i to station j. This ensures that the team at j is completed if it initially lacks firefighters; otherwise, at least one firefighter must relocate to drive the truck.

- **(17)**: The total number of firefighters dispatched from station i cannot exceed its available firefighters c_i.

These personnel rules also apply to the LBAP model. Since LBAP adjusts origin–destination assignments from the MCRP to minimize travel time, any reassignment must respect personnel constraints. If a reassignment would leave a team incomplete, it is disallowed, even if travel time decreases. This ensures relocation plans are realistic and executable.

5.3 Station Dispatch Optimization Model (SDOM)

In this study, we propose the SDOM, a MILP-based approach designed to determine the optimal allocation of vehicles and firefighters at the station level. The model considers vehicle capacity, available personnel, and deployment constraints to maximize dispatch efficiency.

The SDOM focuses exclusively on priority-1 incidents. If the nearest station cannot provide a complete team, other stations are considered in order of increasing distance, and the SDOM is used to compute the optimal dispatch combination.

The objective function follows a hierarchical priority:

1. Maximize the total number of dispatched firefighters
2. Minimize the number of dispatched vehicles
3. Prefer lighter vehicles when the number of vehicles is the same

Parameters and Sets

- $\mathcal{V}$: set of available vehicles at the station
- r: number of firefighters still needed for dispatch
- m: number of firefighters available at the station
- c_i: maximum capacity of vehicle i
- w_i: vehicle preference weight for vehicle i, indicating preference

Decision Variables

- $x_i \in \{0,1\}$: whether vehicle i is dispatched
- $y_i \in \mathbb{Z}_{\geq 0}$: number of firefighters assigned to vehicle i

Objective Function

$$\max \left(\sum_{i \in \mathcal{V}} y_i\right) \cdot \alpha + \left(|\mathcal{V}| - \sum_{i \in \mathcal{V}} x_i\right) \cdot \beta + \sum_{i \in \mathcal{V}} w_i \cdot x_i \tag{18}$$

Constraints

$$\sum_{i \in \mathcal{V}} y_i \leq r \tag{19}$$

$$\sum_{i \in \mathcal{V}} y_i \leq m \tag{20}$$

$$y_i \leq c_i \cdot x_i, \quad \forall i \in \mathcal{V} \tag{21}$$

$$x_i \in \{0,1\}, \quad y_i \in \mathbb{Z}_{\geq 0}, \quad \forall i \in \mathcal{V} \tag{22}$$

Weight Design to Enforce Priority in Objective (18) To guarantee the hierarchical priority in the objective function, we define the weight parameters as follows:

1. The term $(\sum y_i) \cdot \alpha$ aims to maximize the total number of dispatched crew. To ensure that dispatching a single additional firefighter always yields a higher benefit than reducing the maximum possible number of vehicles:

$$\alpha > \beta \cdot |\mathcal{V}|$$

2. The term $(|\mathcal{V}| - \sum x_i) \cdot \beta$ seeks to minimize the number of vehicles. To ensure that reducing one vehicle yields more benefit than fully optimizing vehicle weights:

$$\beta > \sum_{i \in \mathcal{V}} w_i$$

3. The final term $\sum_{i \in \mathcal{V}} w_i \cdot x_i$ introduces a tie-breaker preference for lighter vehicles, applied only when the number of dispatched crew and vehicles are the same.

Here, the vehicle preference weights w_i are assigned based on vehicle type, with representative values:

$$w_{\text{MOTO}} = 0.5, \quad w_{\text{TS2}} = 0.4, \quad w_{\text{TS4}} = 0.3, \quad w_{\text{TS6}} = 0.2$$

To numerically enforce the hierarchy, we define:

$$\beta = \sum_{i \in \mathcal{V}} w_i + 1, \quad \alpha = \beta \cdot |\mathcal{V}| + 1$$

Constraint Interpretations

- **(19)**: Ensures that the total number of dispatched firefighters does not exceed the remaining crew required for the incident.
- **(20)**: Limits the total number of firefighters dispatched to what is available at the current station.
- **(21)**: Ensures the number of firefighters assigned to a vehicle does not exceed its capacity, and firefighters can only be assigned if the vehicle is dispatched.
- **(22)**: Defines the domains of the decision variables: x_i must be binary, and y_i must be a non-negative integer.

6

Simulation

6.1 Simulation Framework

This chapter describes the architecture of our custom-built simulation framework, developed to evaluate the performance of different dispatch strategies, resource configurations, and vehicle compositions under dynamic incident scenarios. The simulator captures the full emergency response process by integrating real incident data, rule-based dispatching, vehicle relocation, and multi-vehicle deployment logic.

6.1.1 Overview of Simulation Design

The framework operates on a minute-by-minute timeline, denoted by t, and replicates the entire response process from incident generation to vehicle recovery or relocation.

Before the main simulation begins, a series of preprocessing steps are performed during initialization, including historical incident sampling, incident duration fitting, arrival rate forecasting, and response time distribution estimation. These steps are executed once and are not shown in the primary workflow.

After preprocessing, each simulation run proceeds through the following stages, as illustrated in Figure 6.1:

- **Incident Generator:** loads and schedules incidents from historical data.
- **Incident Upgrading:** allows incidents to escalate probabilistically over time, possibly requiring additional vehicles.

- **Vehicle Dispatch:** determines which vehicles to assign using a configurable dispatch rule.
- **Dispatch Rule:** pre-defined policies for dispatch. Legacy rule sends fully crewed vehicles; new rule allows partial crewing. The SDOM optimizes allocation for priority 1 incidents under the new rule.
- **Relocation Planner:** relocates idle vehicles to better cover anticipated future demand, based on optimization models.
- **MCRP and LBAP:** two linear programming models used within the relocation planner to identify relocation targets and assign specific vehicles.
- **Vehicle Return:** simulates the recovery and availability of vehicles after deployment or relocation.

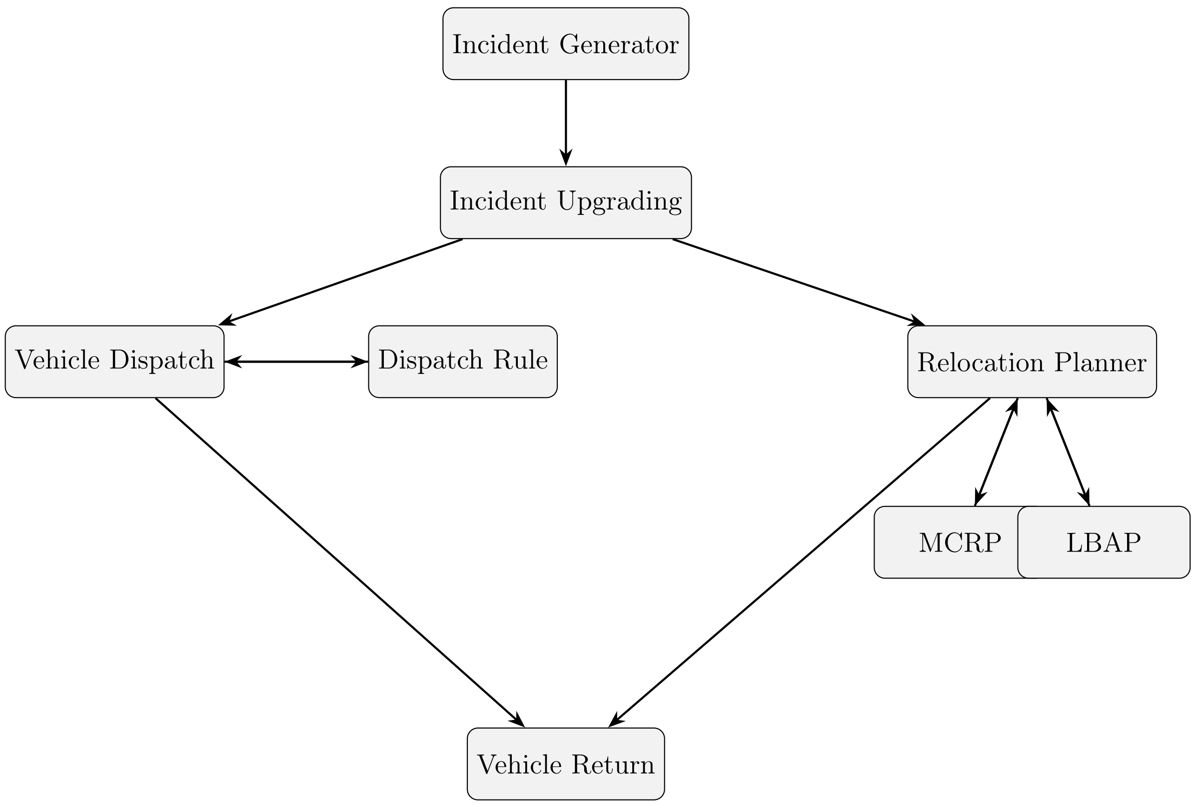

Figure 6.1: Main modules and workflow of the simulation framework, executed after preprocessing

6.2 Incident Generation

In the simulator, incidents are generated following the approach used in the Data Driven Fire Department project (5), focusing on three main aspects: incident type, priority, and incident location.

The simulator stochastically generates incidents for each hour according to the hourly arrival rates estimated by the Prophet forecaster. For each incident, type, priority, and location are sampled according to historical frequencies. High-urgency events, such as indoor fires and resuscitations, are more likely to be generated as priority 1 incidents, while locations with higher historical occurrence rates receive proportionally higher generation probabilities.

In previous versions, the simulator also generated building type to determine target response times, since certain facilities (e.g., prisons, shopping centers) require shorter response times. However, the updated dataset from the fire department system no longer includes building type information. As a result, the simulator now assumes the standard target response time for all incidents. The building type function remains in the code and can be reactivated if such information becomes available in the future.

Once incident type and priority are determined, the simulator selects the required vehicle dispatch. Unlike the previous simulator, which implicitly modeled incident scale by assigning vehicle sets derived from historical records, our framework explicitly defines vehicle requirements for different scale levels. A dispatch rule table, developed with input from domain experts, specifies the vehicle composition for each incident type at each scale.

The dispatch rule table applies to incidents of all priority levels. The numbers indicate the required number of teams—for example, one TS2 represents a two-person team, while two TS4 vehicles correspond to two four-person teams. A new dispatch rule, TS6_P, specifies six personnel without restricting vehicle type, primarily used when indoor fires escalate from small to medium scale. Vehicle requirements for all scale levels are listed in the Appendix, and table 6.1 illustrates the initial (small-scale) vehicle combinations for each incident type.

Incident Type	**RV**	**HV**	**WO**	**TS2**	**TS4**	**TS6**	**TS6_P**	**MOTO**	**DRONE**
General assistance				1					1
Police assistance				1					
Alarm				1					1
Ambulance assistance	1			1					
Fire alarm					1				1
Animal assistance			1		1				1
Outdoor fire					1				1
Resuscitation								1	
Elevator entrapment		1							
Person in water			1		1				1
Hazardous substances				1					
Indoor fire	1				2				1
Extreme weather	1				1				
Other fire					1				1
Vehicle in water		1	1		1				1
Water rescue			1		1				1
Animal in water			1		1				1
Aviation rescue		1			2				1
Unknown				1					

Table 6.1: Vehicle dispatch table for small incident types

6.3 Incident Duration and Upgrading Mechanism

The fitting methods for incident duration and upgrade probabilities were introduced in Chapter 4. Here we describe their implementation in the simulator.

Each incident starts at the small scale, and its final level is drawn from the upgrade probabilities of its type. For types with insufficient data, both upgrade probabilities and duration distributions fall back to the overall fitted parameters.

The simulator generates the total duration together with registration times for possible upgrades. Because the dataset only records the final scale, we designed an *upgrade ratio* to estimate intermediate timings:

- For **major** scale, the upgrade time is sampled as a truncated normal distribution centered at 0.8 (range 0.75–0.85) multiplied by the very major registration time.
- For **medium** scale, the ratio is centered at 0.5 (range 0.45–0.55) relative to the final (**major** or **very major**) registration time.

This ensures consistent timings for all stages. When a stage is reached, the simulator dispatches additional vehicles according to the dispatch rule table.

While each final scale (S, M, J, V, see Table 6.2) has its own fitted duration distribution, these distributions reflect incidents that ended at that scale. They cannot be used to generate intermediate upgrade times for incidents that eventually grow larger—for example, using the distribution of final-medium incidents to determine the medium stage of a final-major incident would mix independent cases and lead to inconsistencies. This could even result in smaller-scale stages lasting longer than larger ones. To avoid such conflicts, we adopt the ratio-based approach, which ensures monotonic scaling of upgrade times. Future work could refine these ratios or, if data permit, use explicit upgrade timestamps.

Algorithm 1 Incident upgrading process

Input: incident type i at time t_0

Output: final level $L \in \{S, M, J, V\}$, end time t_e, registration times (t_M, t_J, t_V)

if $i \in \mathcal{I}_{stats}$ **then** retrieve (p_{SM}, p_{MJ}, p_{JV}) and lognormal parameters $(\theta_S, \theta_M, \theta_J, \theta_V)$;
else use generic (p_{SM}, p_{MJ}, p_{JV}) and $(\tilde{\theta}_S, \tilde{\theta}_M, \tilde{\theta}_J, \tilde{\theta}_V)$;

Sample $r \sim U(0, 1)$ Initialize $t_M = t_J = t_V = \infty$ Sample ratio factors $\rho_J \sim \mathcal{N}_t(0.8, 0.02)$, $\rho_M \sim \mathcal{N}_t(0.5, 0.02)$

if $r < p_{SM} p_{MJ} p_{JV}$ **then**
  $L \leftarrow V$ sample $d_V \sim \text{Lognorm}(\theta_V)$ until $t_V < d_V$ $t_J \leftarrow \rho_J t_V$, $t_M \leftarrow \rho_M t_V$ $t_e \leftarrow t_0 + d_V$
else
  if $r < p_{SM} p_{MJ}$ **then**
    $L \leftarrow J$ sample $d_J \sim \text{Lognorm}(\theta_J)$ until $t_J < d_J$ $t_M \leftarrow \rho_M t_J$ $t_e \leftarrow t_0 + d_J$
  else
    if $r < p_{SM}$ **then**
      $L \leftarrow M$ sample $d_M \sim \text{Lognorm}(\theta_M)$ until $t_M < d_M$ $t_e \leftarrow t_0 + d_M$
    else
      $L \leftarrow S$ sample $d_S \sim \text{Lognorm}(\theta_S)$ $t_e \leftarrow t_0 + d_S$

return (L, t_e, t_M, t_J, t_V)

Table 6.2: Notation used in the incident upgrading process

Symbol	Meaning
S, M, J, V	Incident levels: Small, Medium, Major, Very Major
$\mathcal{I}_{stats}$	Incident types with specific fitted parameters
p_{SM}, p_{MJ}, p_{JV}	Upgrade probabilities between consecutive levels
$\theta_S, \theta_M, \theta_J, \theta_V$	Lognormal parameters for **total duration** by level
$\tilde{\theta}_S, \tilde{\theta}_M, \tilde{\theta}_J, \tilde{\theta}_V$	Generic lognormal parameters (for types not in $\mathcal{I}_{stats}$)
d_S, d_M, d_J, d_V	Total incident duration if final level is S, M, J, V respectively
t_0	Incident starting time
t_e	Incident ending time $t_0 + d$
t_M, t_J, t_V	Registration times for upgraded levels (when the incident is marked as upgraded)
ρ_M, ρ_J	Random scaling factors for registration times, sampled from truncated normal distributions
r	Random variable uniformly sampled from $[0, 1]$ for deciding upgrade path

6.4 Response Timing and Dispatch Design

The simulator models the response process in three phases: response, on-scene, and return. Unlike the earlier implementation, which estimated incident duration by adding fitted on-scene times to the response period, the new design first generates a total incident duration and then derives the response window from it.

6.4.1 Coupling On-Scene Time and End Time

In the previous implementation, the on-scene time was fitted from historical data and incident termination was computed as incident start time + response time + on-scene time.

This approach did not explicitly control the overall duration. In the revised design, an initial end time is drawn for each incident. If vehicles arrive later than the target threshold of 8 minutes, such as when distant stations must respond, the excess delay is evenly distributed across all responding vehicles and added to the incident duration.

Specifically, we assume a target response time of 8 minutes. Any delay beyond this threshold is distributed equally among the dispatched vehicles, increasing the incident duration by (arrival delay – 8 minutes) ÷ number of responding vehicles.

For example, a small indoor fire with a base duration of 15 minutes requires two TS4 units. If their response times are 15 and 20 minutes, the adjusted duration becomes $15 + (15–8)/2 + (20–8)/2 = 24.5$ minutes.

Early arrivals shorten the duration in the same manner. This adjustment affects only the total length of the incident, not the escalation timing, and is intended to prevent unrealistic outcomes in the simulation.

6.4.2 Travel Time Modeling for Different Vehicle Types

Several vehicle types are included: **MOTO**, **TS2**, **TS4**, **TS6**, **BUS**, and **DRONE**. Their travel times are derived from OSRM estimates with type-specific adjustment factors, based on expert judgment:

- TS6 = OSRM × 1.10
- TS4 = OSRM × 1.05
- TS2 = OSRM × 1.00
- MOTO = OSRM × 0.80
- BUS = OSRM × 1.00
- DRONE = OSRM × 0.50

BUS units are used to transport TS crews when no staffed trucks are available; they carry personnel but not equipment and must be dispatched together with fire trucks. DRONEs, included at the request of the fire department, provide rapid reconnaissance. They operate without crew and do not influence truck arrivals. These parameters approximate travel performance and can be refined with further data.

6.5 Dispatch Rules and Multi-Vehicle Assignment

We implement two types of dispatch rules:

1. **Legacy rule** — current operational practice.
   All dispatched vehicles must be fully crewed, regardless of incident type or priority. Even low-demand incidents (e.g., resuscitations) require a full TS6 crew under this rule.

2. **New rule** — flexible deployment with additional vehicle types.
   The new rule differentiates between priority 1 and non-priority 1 incidents:

   - *Priority 1 incidents:* Required vehicles must be included in the dispatch if specified, but additional crew demand may be fulfilled by multiple vehicles of different types. When no required vehicle is immediately available, the simulator allows split dispatch to satisfy personnel requirements.
   - *Non-priority 1 incidents:* Dispatch strictly follows the vehicle table. If the exact type is unavailable but adequate personnel exist, a larger vehicle type may substitute, but multiple vehicles cannot collectively fulfill a single vehicle requirement.

   This flexible scheme reduces unnecessary overstaffing and makes better use of the available fleet. For priority 1 cases, the **SDOM** is applied. Stations are considered in order of increasing travel time:

   (a) Check whether the required vehicle type is available; if so, dispatch it with as many crew as possible.

   (b) If additional personnel are needed, check for available BUS units to supplement crew transport.

   (c) If neither condition is met, use the SDOM to select the best combination of vehicles and personnel from the station.

   (d) Remaining crew requirements are carried over to subsequent stations until fully satisfied.

   This mechanism can also be applied to the current deployment (TS6 only), allowing vehicles to be dispatched with only the required crew (e.g., a TS6 with two firefighters for a TS2-scale incident) and, when necessary, multiple vehicles to jointly meet personnel demands, thus avoiding unnecessary overstaffing (as in Scenario 2 and Scenario 4, Table 7.1).

Pseudocode for Multi-Vehicle Dispatch Logic The following pseudocode summarizes the core logic of the dispatch algorithm used for priority 1 incidents:

Algorithm 2 Vehicle dispatch with multi-vehicle assignment (for $prio = 1$ incidents)

Input: incident location l, required vehicle type v_r, required crew c_r

Output: dispatch list $\mathcal{D} = \{(v_j, \tau_j)\}$ with assigned full-time/part-time crew

Sort all stations $s \in \mathcal{S}$ by travel time $\tau_s(l)$ ascending Initialize remaining crew $c \leftarrow c_r$, required_sent $\leftarrow$ false, $\mathcal{D} \leftarrow \emptyset$ Define TS-series types $\mathcal{T} = \{$TS6, TS4, TS2, MOTO$\}$

foreach $(s, \tau_s(l))$ *in sorted stations* **do**
- **if** $c \leq 0$ **then break**;
- Get available vehicles at s: $\mathcal{V}_s \leftarrow \{v \mid v \text{ free at } s\}$
- `// Step 1: dispatch required vehicle if not yet sent`
- **if** ***not*** *required_sent* **then**
  - **foreach** $v \in \mathcal{V}_s$ **do**
    - **if** *v.type* $= v_r$ ***or*** *(size(v)* $> c_r$ ***and*** *v.type* $\in \mathcal{T}$*)* **then**
      - assign $d \leftarrow \min(c_r, crew\_avail(v), c)$ **if** $d > 0$ **then**
        - set $v.d\_ft \leftarrow \min(d, ft_s)$, $v.d\_pt \leftarrow d - v.d\_ft$ append $(v, adj\_travel(v, \tau_s(l)))$ to $\mathcal{D}$ $c \leftarrow c - d$, required_sent $\leftarrow$ true **break**
- **if** $c \leq 0$ **then break**;
- `// Step 2: check BUS dispatch`
- **foreach** $v \in \mathcal{V}_s$ **do**
  - **if** *v.type* $= BUS$ **then**
    - assign $d \leftarrow \min(crew\_avail(v), c$ **if required_sent else** $c - 1)$ **if** $d > 0$ **then**
      - set $v.d\_ft \leftarrow \min(d, ft_s)$, $v.d\_pt \leftarrow d - v.d\_ft$ append $(v, 1.1\tau_s(l))$ to $\mathcal{D}$ $c \leftarrow c - d$ **break**
- **if** $c \leq 0$ **then break**;
- **if** ***BUS dispatched at*** s **then**
  - **continue**
- **if** ***not*** *required_sent* ***and*** $c = 1$ **then**
  - **continue**
- `// Step 3: SDOM optimization for remaining TS vehicles`
- $\mathcal{V}_s^{TS} \leftarrow \{v \in \mathcal{V}_s \mid v.type \in \mathcal{T}\}$ $(\hat{\mathcal{V}}, crew\_alloc) \leftarrow SDOM\_optimize(\mathcal{V}_s^{TS}, c, ft_s, pt_s)$ increment SDOM counter **foreach** $v \in \hat{\mathcal{V}}$ **do**
  - assign adjusted travel time $\tau_v \leftarrow adj\_travel(v, \tau_s(l))$ append (v, τ_v) to $\mathcal{D}$ $c \leftarrow c - crew\_alloc(v)$

if $c > 0$ **then**
- append (None, None) to $\mathcal{D}$; `// external vehicle required`

return $\mathcal{D}$

Table 6.3: Notation used in the vehicle dispatch algorithm

Symbol	Meaning
l	Incident location
v_r	Required vehicle type (e.g., TS6)
c_r	Crew required for v_r
$\mathcal{S}$	Set of stations
$\tau_s(l)$	OSRM travel time from station s to location l
$\mathcal{T}$	TS vehicle series: TS6, TS4, TS2, MOTO
$size(v)$	Full crew required to operate vehicle v (vehicle capacity)
$crew\_avail(v)$	Available staff at station to crew vehicle v
d	Number of crew to be dispatched on vehicle v
$v.d\_ft, v.d\_pt$	Full-time / part-time staff assigned to vehicle v for this dispatch
ft_s, pt_s	Full-time and part-time staff counts at station s
$adj\_travel(v, \tau)$	Adjusted travel time factor: ×1.1 for TS6, ×1.05 for TS4, ×1 for TS2/BUS, ×0.8 for MOTO
$\mathcal{D}$	Dispatch result list $(vehicle, travel\_time)$
SDOM_optimize	Optimization subroutine to select vehicles/crew at a station
(None, None)	Indicates unmet crew requirement after all stations are checked (external vehicle needed)

6.6 Vehicle Relocation and Return

6.6.1 Relocation and Return Design

When a major or very major indoor fire occurs, a relocation algorithm is triggered to redistribute vehicles across stations. The aim is to prevent local resource shortages and maintain adequate coverage in unaffected areas.

Relocated vehicles are assigned a `support_until_t` parameter, initially set to the estimated end time of the triggering incident. If this estimate is updated, the support period is adjusted accordingly. During the support period, relocated vehicles may be dispatched to nearby incidents. After completing such calls, they first return to the temporary station they were relocated to, rather than directly to their base station.

Once the support period ends and the vehicle becomes available, it must return to its base station. For consistency in simulator logic, the vehicle always returns via its temporary station before heading back, rather than traveling directly from an incident scene to its home base.

6.6.2 Trigger Mechanism

The relocation algorithm is currently activated under the following rule:

- Any major or very major indoor fire triggers relocation.
- A very major fire triggers relocation twice: once upon escalation to major, and once again when it reaches very major.

This is a simplified setup, intended as a baseline for further development. More realistic triggers could include coverage thresholds (e.g., relocation only when effective coverage falls below a predefined level) or concurrency rules (e.g., when more than three indoor fires occur simultaneously). Exploring such extensions is left for future work.

6.6.3 Relocation Algorithm

To enhance system resilience and dynamically rebalance coverage, relocation plans are generated using two optimization models: the MCRP and the LBAP. The simulator applies these plans by updating vehicle positions and assigning the `support_until_t` parameter, which limits how long relocated vehicles can be dispatched outside their home base. During this period, they are treated the same as vehicles permanently stationed there.

6.6.4 Required Input Data

- **Travel times:** Pairwise travel times between fire stations and demand locations, derived from OSRM.
- **Current deployments:** The current distribution of vehicles among stations.
  - Only vehicles currently located at their *base stations* are considered eligible for relocation. This avoids selecting vehicles that have already been relocated and not yet returned, which would introduce significant complexity into both simulation logic and operational feasibility.
- **Demand rates:** Estimated spatial probabilities of incident occurrence, aggregated per station based on historical frequency at associated demand locations.

- **Response neighbourhood matrix (A):** A binary matrix where $a_{ik} = 1$ indicates that station $i \in N$ can serve neighbourhood $k \in K_n$, based on travel-time proximity.

  In this study, we set $n = 3$, meaning each neighbourhood can be served by its three nearest stations. This configuration results in 44 response neighbourhoods across the study area (see Figure 6.2). The simulator also allows users to specify different values of n, which would yield alternative response neighbourhood configurations.

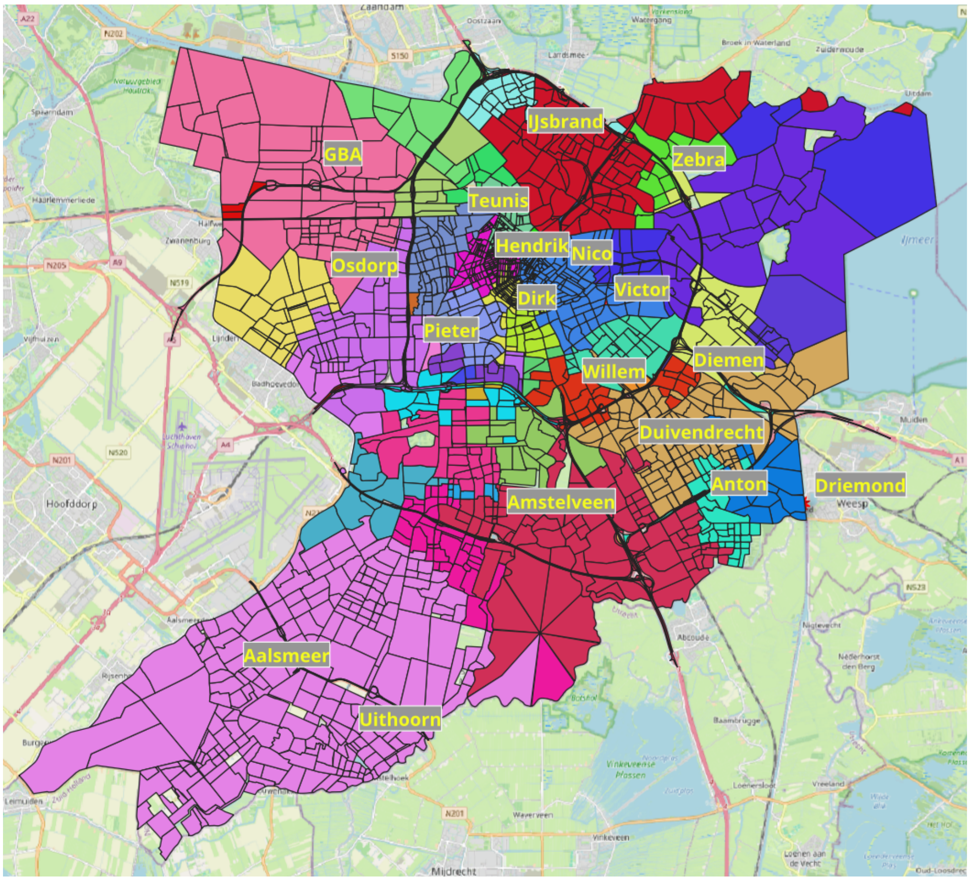

Figure 6.2: Response neighbourhoods ($n = 3$) derived from OSRM travel times

6.7 Summary and Future Improvements

The proposed simulation framework offers a flexible and modular environment to evaluate emergency response strategies under realistic conditions. Its main strengths include:

- A dynamic incident upgrading mechanism based on fitted distributions.
- Support for multi-vehicle dispatch and flexible vehicle substitution logic.

- Integration of relocation strategies using mathematical optimization (MCRP + LBAP).
- Customization of dispatch rules, including the SDOM.
- Differentiation between multiple vehicle types and dynamic availability.

However, several limitations remain, suggesting directions for future development:

- **Coverage-based Relocation Trigger:** current relocation is periodic and could be enhanced with real-time coverage evaluation to trigger relocations only when necessary.
- **More Realistic Upgrading Paths:** incident escalation processes could be modeled using agent-based logic or Markov chains to simulate resource intensification more realistically.
- **Richer Vehicle Types & Traffic Inputs:** extending the model to include additional vehicle types (e.g., ladder trucks, ambulances) and dynamic travel times based on traffic data would improve realism.

7

Experiment

7.1 Experimental Setup

To evaluate the performance of different vehicle deployment strategies and dispatch operations, we designed a set of simulation scenarios based on realistic configurations in the Amsterdam-Amstelland fire region. The main factors of variation are vehicle deployment (legacy vs. new configurations), dispatch rules (current vs. redesigned), and the use of a relocation strategy.

Vehicle Deployment Scenarios

We distinguish between two main vehicle deployment settings. The crew numbers shown here represent the staffing required to operate the deployed vehicles, not the total number of firefighters employed by the department. In reality, to sustain year-round operations, the fire department employs several times more personnel than those directly assigned to vehicles.

- **Legacy Deployment:** This reflects the current operational setup used by the fire department. Only TS6 vehicles are deployed, with staffing patterns based on the existing full-time and part-time crew assignments. The total number of TS6 vehicles is 22, supported by 84 full-time and 48 part-time firefighters.
- **New Deployment:** In collaboration with domain experts, a new configuration was designed incorporating additional vehicle types—namely MOTO, TS2, and TS4—while reducing overall crew demand. The new scenario deploys 18 MOTO, 10 TS2, 12 TS4, and 9 TS6 vehicles across the region, supported by 73 full-time and 28 part-time firefighters.

Each station is equipped with one BUS and one DRONE under both deployments. Tables A.4 and A.5 provide a detailed overview of the station-level allocations. While the new configuration implies additional procurement and operational costs, these are not explicitly modelled in our analysis. Instead, we use the expert-proposed setup as a reference for operational evaluation. Future work could extend the model by including cost-efficiency considerations.

Although the simulator supports a wide range of vehicle types (e.g., RV, HV, DRONE), the analysis in this chapter focuses on the TS-series vehicles, which are central to urban firefighting operations. Results for other vehicle types are produced but not discussed in detail.

All scenarios are simulated over a “100-year”. The “100-year” simulation does not represent a continuous century of operations. Rather, it consists of 100 independent one-year runs, each with its own randomly generated incident stream. At the end of each year, vehicle states are reset and incidents restart from day one. Final results are aggregated across all runs to obtain stable estimates.

This design was chosen mainly due to computational limitations of the platform, which was run on a machine with an Intel(R) Core(TM) Ultra 5 125U processor (1.3 GHz, 12 cores, 14 logical processors). Running a single continuous 100-year simulation with full incident and vehicle history was not feasible. Moreover, the year-by-year reset mitigates the effect of early random fluctuations—for example, if a single station happens to be overloaded at the start of the simulation, this could trigger cascading delays. The reset provides a neutral starting point each year and improves statistical fairness across runs.

Scenario Design

For both deployment configurations, we simulate four operational scenarios, resulting in a total of eight test cases. The scenarios differ in dispatch rules and whether a relocation strategy is applied:

Each scenario is simulated over a 100-year equivalent of incident data. To assess system performance under stress, all scenarios are also evaluated under high-demand conditions, where incident arrival rates are doubled to emulate extreme operational pressure — a situation of particular interest to the fire department.

Evaluation Metrics

The primary performance indicator in this study is **response time**, defined as the interval between the incident start time and the arrival time of the first vehicle. It reflects op-

Table 7.1: Operational scenarios for simulation experiments

Scenario	Vehicle Deployment	Dispatch Rule	Relocation Strategy
1 (benchmark — current system)	Old deployment	Legacy	No
2	Old deployment	New	No
3	Old deployment	Legacy	Yes
4	Old deployment	New	Yes
5	New deployment	Legacy	No
6	New deployment	New	No
7	New deployment	Legacy	Yes
8	New deployment	New	Yes

erational effectiveness and compliance with legal response-time targets. We compute the proportion of incidents responded to within the legally defined thresholds, which vary by location and risk level.

We define the *Fraction of On-Time Arrivals* (FOTA) at threshold T as the share of incidents with response times not exceeding T. A higher FOTA indicates better coverage and faster intervention.

We also track the rate of incidents requiring **external assistance**, i.e., cases where no TS-series vehicle is available within the region, prompting support from neighboring service areas. This is especially relevant for high-priority incidents, such as resuscitations or indoor fires, where delayed responses can have serious consequences.

To gain further insight, metrics are analyzed not only for all incidents but also for selected high-priority types:

- Indoor fires ("Binnenbrand")
- Resuscitations ("Reanimatie")
- Periods of major incidents: defined as times when an ongoing indoor fire is categorized as "major" or "very major"

The last category is particularly important, as the relocation strategy is only triggered during large-scale incidents. While other incident types could also require relocation in practice, we limit activation to indoor fires in this study to keep the scenarios clear and controlled. For analyses covering all incidents and indoor fires specifically, an 8-minute threshold is used; for resuscitations, a 6-minute threshold is applied, consistent with Stieglis et al. (2025 (17)), which highlights that first-shock delivery within six minutes is a key survival indicator.

7.2 Results

This section presents the simulation results for the eight scenarios described in Section 7.1. We first report the outcomes for the normal arrival-rate experiments (using the hourly arrival rates estimated by the Prophet forecaster), followed by the high-demand experiments (arrival rates scaled by 2), and finally the results for *Big incident* periods (periods with at least one major or very-major indoor fire). For each set of experiments, we provide: (i) relocation-algorithm activity, (ii) overall statistics for all incidents, (iii) indoor fire ("Binnenbrand") results, and (iv) resuscitation ("Reanimatie") results. Where relevant, we also report station-level FOTA (Fraction Of Time Available) rankings and statistical tests. All detailed per-scenario external-support tables are collected in Appendix A.4.

7.2.1 Normal Arrival Rate (Baseline)

7.2.1.1 Relocation Algorithm Execution

Table-like summary of relocation algorithm activations and executions in the relocation-enabled scenarios under normal arrival rates:

- **Old deployment scenarios (TS6 only)**
  - Scenario 3: `Relocation_activated = 2720`, `Relocation_performed = 2662 (97.87%)`
  - Scenario 4: `Relocation_activated = 2795`, `Relocation_performed = 2763 (98.86%)`
- **New deployment scenarios (new vehicles)**
  - Scenario 7: `Relocation_activated = 2812`, `Relocation_performed = 2071 (73.65%)`
  - Scenario 8: `Relocation_activated = 2793`, `Relocation_performed = 969 (34.69%)`

Interpretation: The number of relocation triggers is similar across old and new deployment scenarios, indicating that the occurrence of major indoor fires is comparable. In the old deployment, relocations were almost always executed once triggered. In the new deployment, however, many triggered relocations could not be carried out, particularly in Scenario 8. This is due to the smaller number of TS6 vehicles available (the current implementation only relocates TS6s) and to the new dispatch rule, which assigns TS6s more frequently to smaller incidents, increasing their utilization and reducing their availability for relocation.

7.2.1.2 Aggregate Behaviour (All Incidents)

When aggregating all incident types, the overall fraction of cases requiring external assistance remains below 0.01% across all scenarios (complete counts are provided in Appendix A.4, Table A.4.0.1).

Station-level FOTA and ranking. We calculated FOTA for each station and ranked the stations according to their performance.

Table 7.2: Station-level FOTA (All periods, All incidents, normal arrival rate)

station_letter	Scenario_1	Scenario_2	Scenario_3	Scenario_4	Scenario_5	Scenario_6	Scenario_7	Scenario_8
A	49.70%	56.33%	49.27%	56.94%	60.61%	56.88%	60.52%	57.18%
B	48.49%	53.80%	48.67%	54.33%	59.65%	61.76%	59.17%	61.88%
C	50.29%	55.85%	50.11%	56.47%	60.08%	62.46%	60.27%	62.42%
D	50.40%	56.21%	50.42%	56.38%	60.50%	59.30%	60.34%	59.43%
E	58.86%	63.64%	59.02%	64.07%	69.51%	66.95%	69.05%	66.94%
F	57.52%	62.46%	57.45%	63.06%	66.16%	65.28%	65.60%	64.63%
G	35.94%	40.18%	35.59%	39.83%	42.94%	42.90%	43.10%	43.32%
H	46.31%	51.93%	46.72%	52.44%	55.58%	56.30%	55.78%	56.36%
I	52.92%	58.84%	52.89%	58.73%	59.03%	61.13%	59.35%	61.41%
J	24.04%	27.36%	24.41%	26.76%	30.39%	28.43%	30.99%	28.65%
K	46.89%	51.43%	46.67%	51.94%	54.65%	53.42%	54.94%	52.92%
L	57.05%	61.93%	57.37%	61.63%	59.79%	62.67%	59.39%	62.24%
M	49.02%	55.04%	48.62%	54.84%	58.15%	61.35%	58.00%	60.79%
N	66.03%	69.44%	65.90%	70.15%	76.23%	76.13%	76.59%	75.96%
O	36.93%	40.05%	36.08%	40.75%	41.94%	41.18%	41.64%	40.93%
P	47.40%	53.40%	47.71%	54.13%	63.37%	62.12%	63.04%	61.97%
Q	29.73%	32.99%	29.56%	33.26%	33.59%	35.60%	34.06%	36.01%
R	51.82%	57.50%	52.05%	57.61%	58.91%	58.60%	58.59%	58.39%

Tables 7.2 and 7.3 summarize the station-level FOTA values and their rankings. Figure 7.1 illustrates FOTA curves for all stations and scenarios.

Overall, scenarios with the new vehicle mix (Scenarios5–8) achieve higher FOTA compared to the original deployment (Scenarios1–4), indicating shorter response times. Variation among the new-deployment scenarios is minimal.

Scenario 1 represents the current operational model (only TS6) without the new dispatch rules or relocation, whereas Scenario 5 introduces the new vehicle types while keeping the operational mode otherwise unchanged. The results demonstrate that incorporating the new vehicle mix significantly enhances performance: overall FOTA increases from 33.64% to 47.50%, an absolute improvement of 13.86 percentage points, corresponding to a relative gain of approximately 41.20%.

Given that the response time distributions are visibly non-normal (Figure 7.2), potentially log-normal, we applied the Mann–Whitney U test to compare medians.

Table 7.3: Station FOTA ranks (All periods, All incidents, normal arrival rate)

station_letter	Scenario_1	Scenario_2	Scenario_3	Scenario_4	Scenario_5	Scenario_6	Scenario_7	Scenario_8
A	7	6	8	4	1	5	2	3
B	8	6	7	5	3	2	4	1
C	7	6	8	5	4	1	3	2
D	8	6	7	5	1	4	2	3
E	8	6	7	5	1	3	2	4
F	7	6	8	5	1	3	2	4
G	7	5	8	6	3	4	2	1
H	8	6	7	5	4	2	3	1
I	7	5	8	6	4	2	3	1
J	8	5	7	6	2	4	1	3
K	7	6	8	5	2	3	1	4
L	8	3	7	4	5	1	6	2
M	7	5	8	6	3	1	4	2
N	7	6	8	5	2	3	1	4
O	7	6	8	5	1	3	2	4
P	7	6	8	5	3	2	4	1
Q	7	6	8	5	4	2	3	1
R	8	6	7	5	1	2	3	4
Mean	7.44	5.61	7.56	5.11	2.39	2.67	2.56	2.67
Rank	7	6	8	5	1	3	2	3

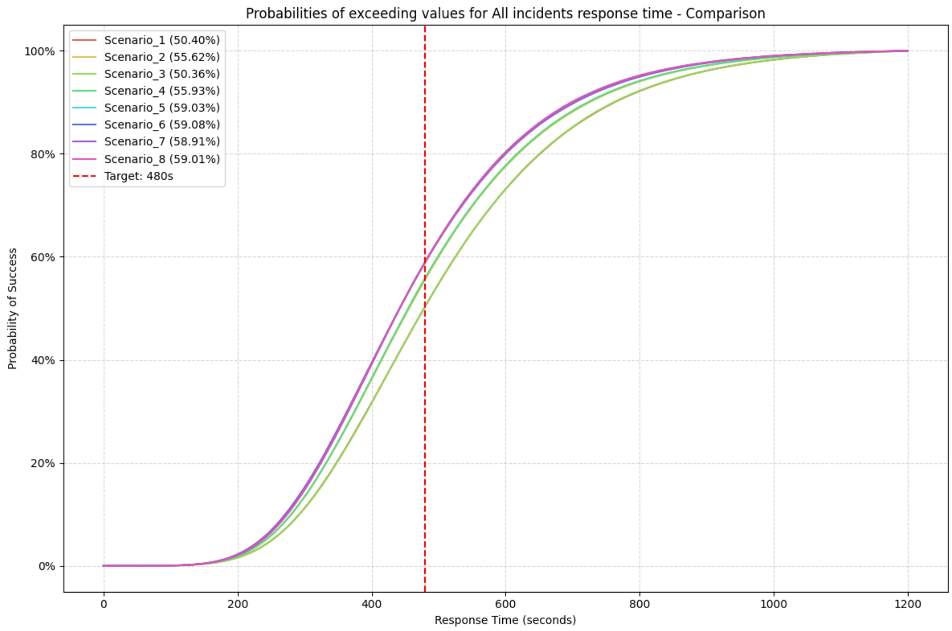

Figure 7.1: FOTA curves for all stations and scenarios (normal arrival rate, all incidents)

The hypotheses were:

- H_0: median_Scenario 1 $\leq$ median_Scenario 5,
- H_1: median_Scenario 1 $>$ median_Scenario 5.

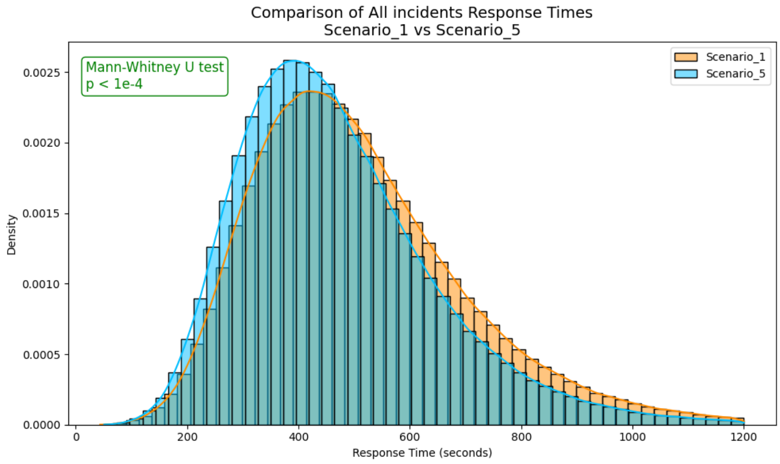

Figure 7.2: Comparison of Scenario 1 (old deployment, legacy dispatch rule) vs Scenario 5 (new deployment, legacy dispatch rule) under normal arrival rate.

The test yields $p < 10^{-4}$; hence, H_0 is rejected. This confirms that Scenario 1 has a significantly higher median response time than Scenario 5, demonstrating that introducing the new vehicle types can improve fire department performance in terms of response time.

Effect of dispatch rule and relocation (All incidents). Figure 7.3 compares response time distributions between the old and new dispatch rules under two deployment settings: (a) Scenario 1 vs. Scenario 2 (no relocation), and (b) Scenario 3 vs. Scenario 4 (with relocation). In both cases, adopting the new dispatch rule shifts the response time distribution towards shorter values, yielding a FOTA improvement of approximately 11–13%.

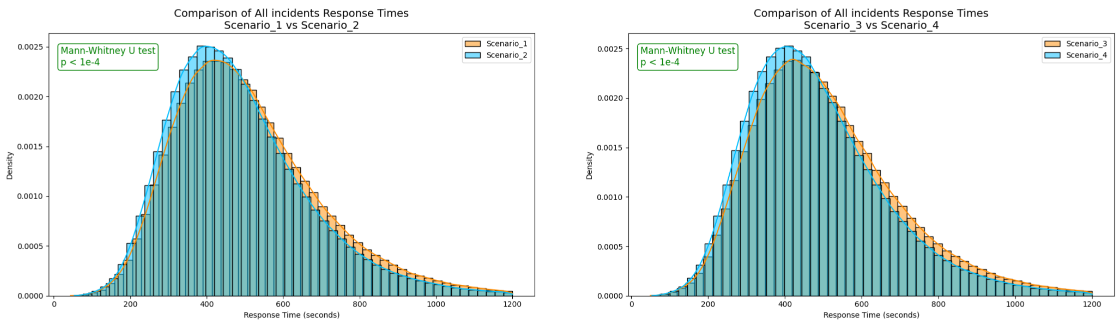

(a) Scenario 1 vs Scenario 2

(b) Scenario 3 vs Scenario 4

Figure 7.3: Impact of switching from old dispatch rule to new dispatch rule under normal arrival rate.

The relocation-only scenario (Scenario 3) shows little improvement over the baseline (Scenario 1) and is actually slightly worse in our simulation results. This is mainly because relocation events are rare compared to the total number of dispatches: out of 865,343 dispatches, the relocation algorithm was triggered only 2,662 times (about 0.31%). Since relocation occurs only for major or very major indoor fires, its overall impact on aggregate performance is limited under the current triggering rules.

7.2.1.3 Indoor Fire (Binnenbrand)

In the simulator, indoor fires have specific vehicle requirements: the initial small level requires two TS4 units, while all higher severity levels require TS6. Consequently, stations without TS4 or TS6 cannot respond to these incidents under the old dispatch rule. In the current deployment, Stations 5, 9, and 17 fall into this category. Use of external support for indoor fires is minimal: across all scenarios, fewer than 0.01% of incidents required external assistance (full counts are in Appendix A.4, Table A.4.0.2).

In the new deployment scenarios, MOTO, TS2, TS4, and TS6 units were redistributed across stations (see Appendix A); not all stations have every vehicle type. The relocation algorithm currently applies only to TS6 units. Under the old dispatch rule (Scenarios 5 and 7), stations without TS4/TS6 could not dispatch for indoor fires. The new dispatch rule, however, allows smaller units (MOTO, TS2) to cooperate across stations, enabling previously limited stations (e.g., 5, 9, 17) to participate in indoor-fire responses in Scenarios 6 and 8.

Figure 7.4 shows FOTA curves for indoor fires under all scenarios. The new dispatch rule (Scenarios 2, 4, 6, 8) clearly improves performance compared to the old rule (Scenarios 1, 3, 5, 7). In the old-deployment setting (Scenarios 1 vs 2, 3 vs 4), the improvement is roughly 28–30%; in the new-deployment setting (Scenarios 5 vs 6, 7 vs 8), the gain is about 15–17%. These differences are supported by the pairwise scenario tests in Figure 7.5, which consistently show an upward shift in the FOTA curve when the new dispatch rule is used.

Relocation under the current triggering rules performs poorly in some comparisons (e.g., Scenario 3 vs 1, Scenario 7 vs 5). This is likely due to overly frequent triggering: the algorithm activates immediately when a major or very-major indoor fire is detected and executes relocation if an optimal solution is found, without further evaluation. In practice, a commander would likely review the plan before execution. Frequent vehicle movements and returns in the simulator can inadvertently increase overall response times. This does not mean the algorithm is ineffective; rather, its triggering and usage policy need careful tuning. Further investigation into adaptive triggers and decision policies is recommended.

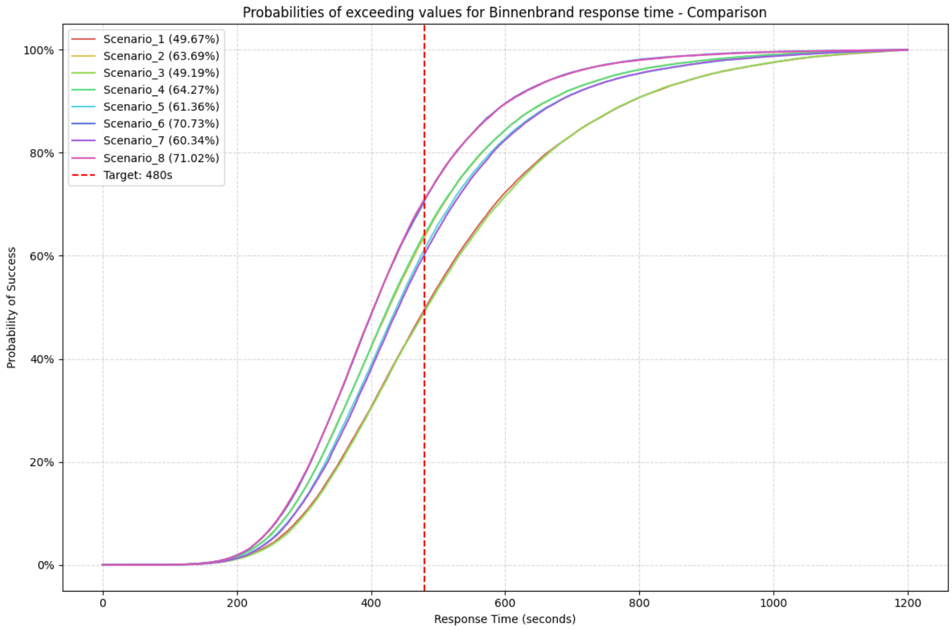

Figure 7.4: FOTA curves for Indoor fire across all scenarios (normal arrival rate).

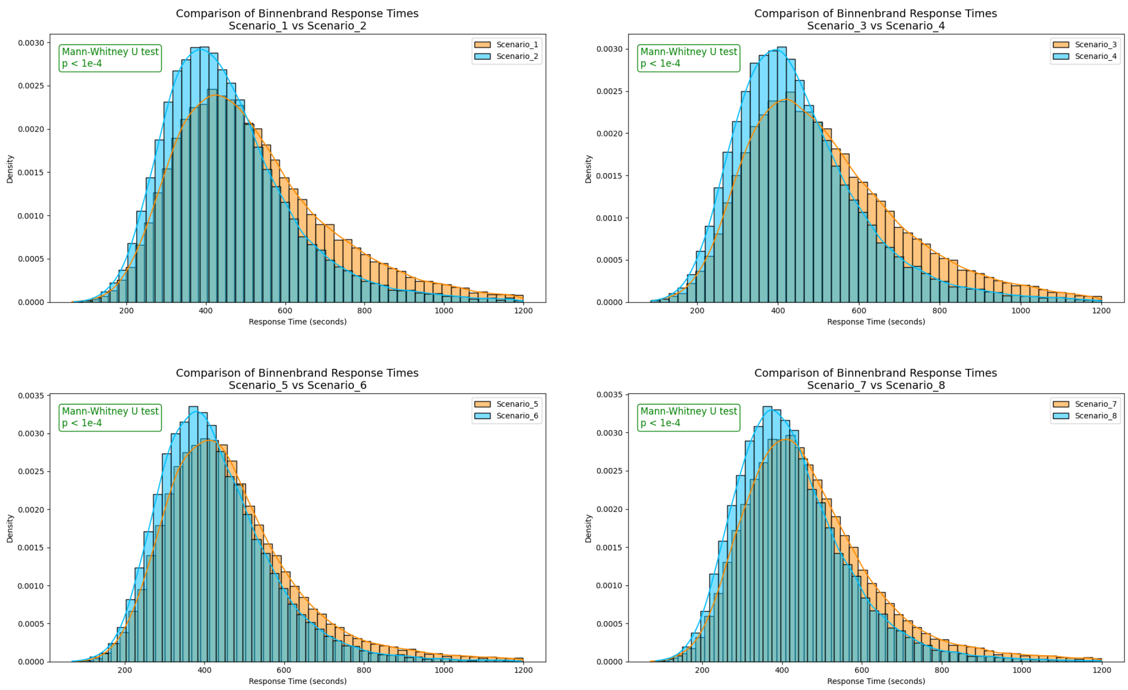

Figure 7.5: Pairwise scenario comparisons for Indoor fire at normal arrival rate.

7.2.1.4 Resuscitation (Reanimatie)

Resuscitation incidents require rapid intervention. In this study, we assume that a single responder carrying an AED is sufficient to handle such incidents, meaning that a TS6 with only one crew member may respond under the new dispatch rule. MOTO or other small units can also assist by transporting personnel or providing support. External support is negligible for this category at normal arrival rates (see Appendix A.4, Table A.4.0.3).

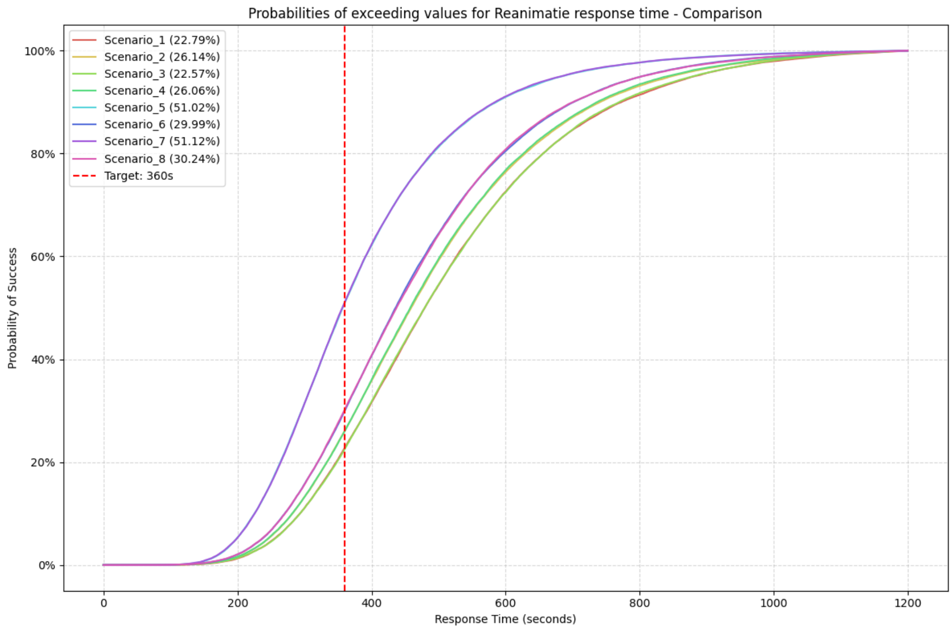

Figure 7.6: FOTA curves for Resuscitation across all scenarios (normal arrival rate).

Figure 7.6 shows that, under the old deployment (Scenarios 1–4), the new dispatch rule improves performance for resuscitation. Allowing single-crew TS6 dispatch increases FOTA by about 15% (Scenario 2 vs 1 and Scenario 4 vs 3). These gains are confirmed by the pairwise statistical tests in Figure 7.7.

In the new deployment (Scenarios 5–8), however, the new rule can slightly reduce resuscitation performance (Scenario 6 vs 5, Scenario 8 vs 7). The main reason is MOTO use: in our configuration, each station has only one MOTO. Under the old rule, MOTOs were used almost exclusively for resuscitation, ensuring availability. The new rule allows MOTOs to support other incidents, which lowers their availability for resuscitation and can reduce performance.

Nevertheless, comparing the deployments without changing the dispatch rule shows that the new vehicle mix itself provides a substantial benefit: FOTA for resuscitation rises from 22.8% in Scenario 1 to 51.0% in Scenario 5, an increase of roughly 124%, highlighting

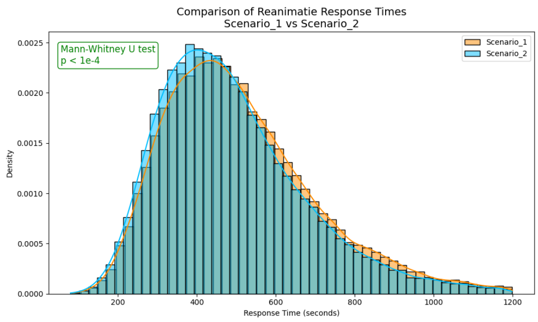

(a) Scenario 1 vs Scenario 2

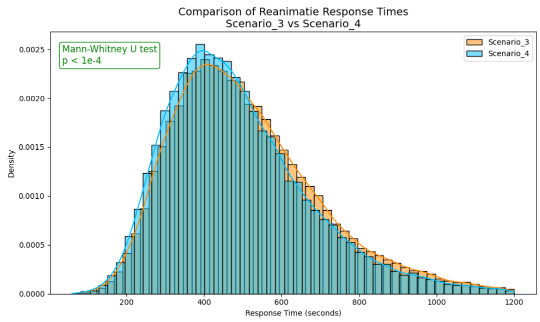

(b) Scenario 3 vs Scenario 4

Figure 7.7: Pairwise comparisons of old deployment scenarios under normal arrival rate (Resuscitation).

the strong positive effect of the updated station-level vehicle allocation.

From an operational view, if new vehicle types are introduced, it may be worth considering whether MOTOs should remain dedicated to resuscitation, rather than being dispatched more broadly.

7.2.2 High-Demand Experiments (Arrival Rates $\times 2$)

To model extreme operational stress we double the incident arrival rates (i.e., the rates estimated by the Prophet forecaster are multiplied by 2). The resulting event counts are approximately twice the normal case.

7.2.2.1 Relocation Algorithm Execution (Scaled Arrival)

- **Old deployment scenarios (TS6 only)**
  - Scenario 3: `Relocation_activated = 5427`, `Relocation_performed = 4882 (89.96%)`
  - Scenario 4: `Relocation_activated = 5462`, `Relocation_performed = 5128 (93.89%)`
- **New deployment scenarios (new vehicles)**
  - Scenario 7: `Relocation_activated = 5665`, `Relocation_performed = 3524 (62.21%)`
  - Scenario 8: `Relocation_activated = 5575`, `Relocation_performed = 1048 (18.80%)`

Because the arrival rates are doubled, the number of incidents also increases by roughly a factor of two. The relocation patterns, however, remain broadly similar to the normal-rate case. Relocation is more readily executed in the old-deployment scenarios, whereas under higher demand and a fixed fleet size, the number of situations without available vehicles increases, lowering the share of successful relocations. This effect is most pronounced in the new-deployment scenarios, particularly in Scenario 8, where fewer than 20% of relocation triggers could be carried out.

7.2.2.2 All Periods: All Incidents (Scaled Arrival)

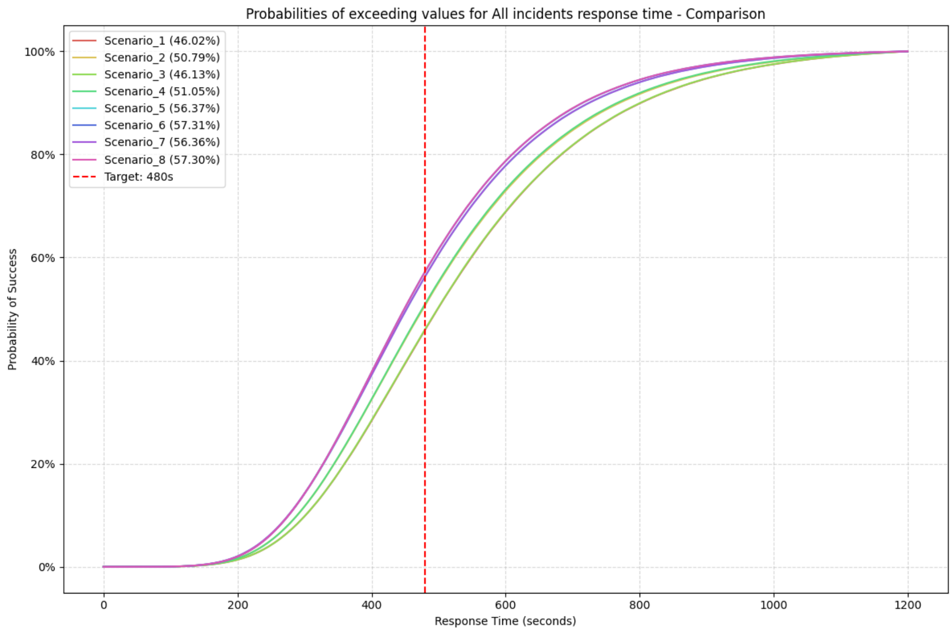

Figure 7.8: FOTA comparisons for all incidents under scaled arrival scenarios.

External support for All incidents is minimal: across all scaled-arrival scenarios, fewer than 0.01% of incidents required external assistance. These values are aggregated over the entire incident set; for rare incident types the proportion may be higher (see Appendix A.4, Table A.4.0.4).

FOTA comparisons for all incidents under scaled-arrival scenarios are presented in Figure 7.8. Comparing the scaled-arrival results with the baseline (normal arrival), FOTA performance in the old-deployment scenarios (1–4) decreases by approximately 8–9% (absolute drop), whereas in the new-deployment scenarios (5–8) the decrease is only about 3–4.5%. In other words, the new-deployment variants show greater resilience to increased demand.

Relative to the baseline (normal arrival), FOTA performance in the old-deployment scenarios (1–4) declines by about 8–9% , while in the new-deployment scenarios (5–8) the decline is limited to –4.5%. The new deployment is therefore more resilient to increased demand.

Within the old-deployment set, the current operational method (Scenario 1) performs worst. Introducing the new dispatch rule produces marked gains, as shown by the pairwise comparisons Scenario 2 vs Scenario 1 and Scenario 4 vs Scenario 3 (Figure 7.9).

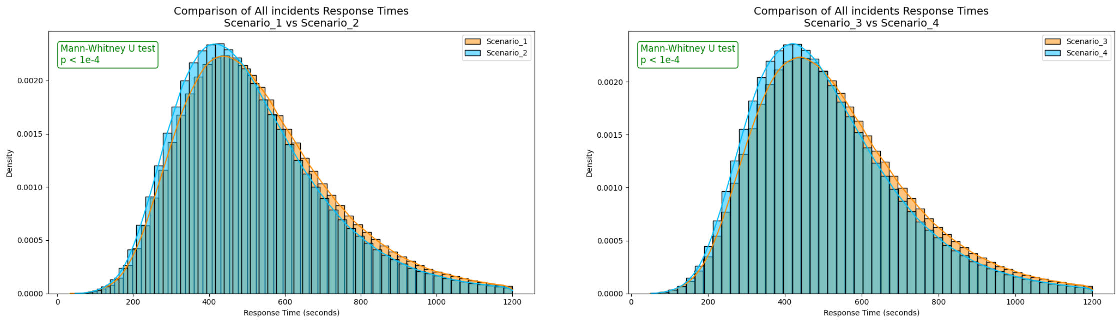

(a) Scenario 1 vs Scenario 2

(b) Scenario 3 vs Scenario 4

Figure 7.9: Pairwise comparisons of old-deployment scenarios under scaled arrival for all incidents.

Even with roughly twice as many incidents, FOTA values in the new-deployment scenarios still exceed those of the old deployment under normal arrival. For instance, Figure 7.10 shows that Scenario 5 (new deployment, scaled arrival) outperforms Scenario 1 (old deployment, normal arrival), underscoring the greater robustness of the new deployment.

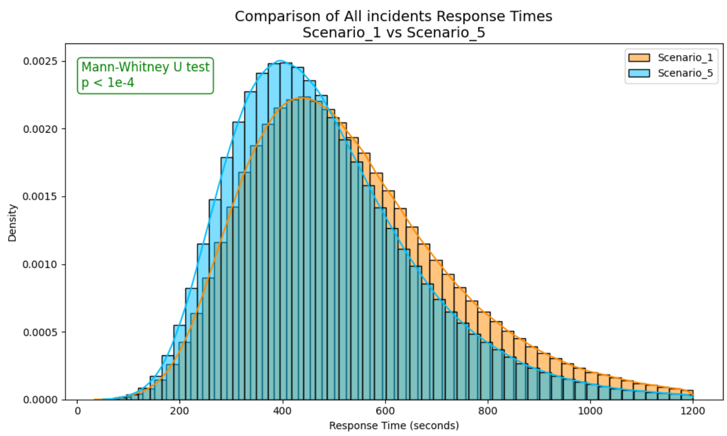

Figure 7.10: Comparison of Scenario 1 (normal arrival, old deployment) vs Scenario 5 (scaled arrival, new deployment) for all incidents.

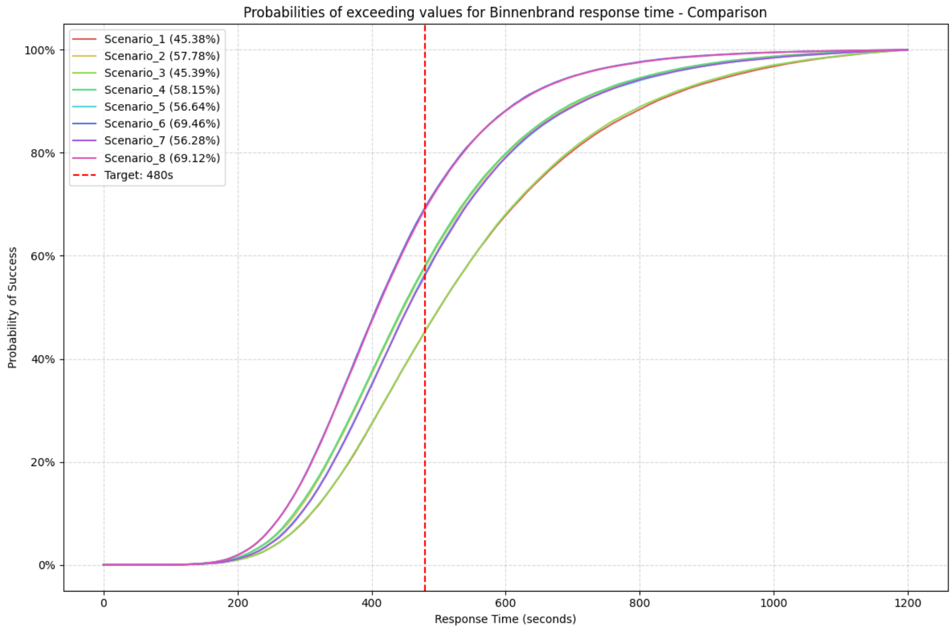

Figure 7.11: FOTA comparisons for Indoor fire under scaled arrival scenarios.

7.2.2.3 Indoor Fire (Scaled Arrival)

External support for Indoor fires is very limited under scaled arrival: across all scenarios, no more than 0.05% of incidents required external assistance. Absolute counts rise compared with the normal-arrival case, but the proportions remain negligible (see Appendix A.4, Table A.4.0.5).

Relative to the normal-arrival experiments, Indoor-fire FOTA in the old-deployment scenarios falls by about 8–10%. In the new-deployment scenarios the decline is smaller: roughly 2% under the new dispatch rule and 7–8% under the old rule. The new dispatch rule therefore provides a consistent gain for Indoor fires.

Figures 7.11 and 7.12 show that the new dispatch rule improves Indoor-fire FOTA across both deployment settings, with the largest gains in the new-deployment scenarios.

7.2.2.4 Resuscitation (Scaled Arrival)

External support for Resuscitation is virtually absent under scaled arrival. In the new-deployment scenarios (5–8), no external-support events are observed, while in the old-deployment scenarios (1–4) the share of incidents requiring external assistance does not exceed 0.01%. Under normal arrival rates, Resuscitation incidents never required external help; the small number of cases appearing under higher demand suggests that, when

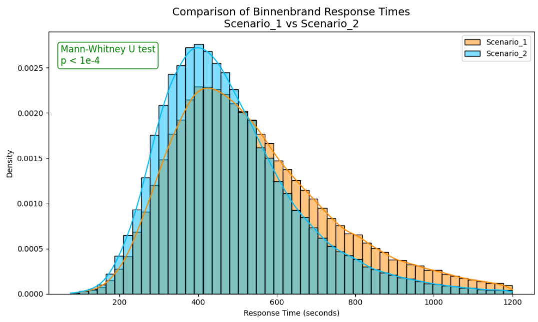

(a) Scenario 1 vs Scenario 2

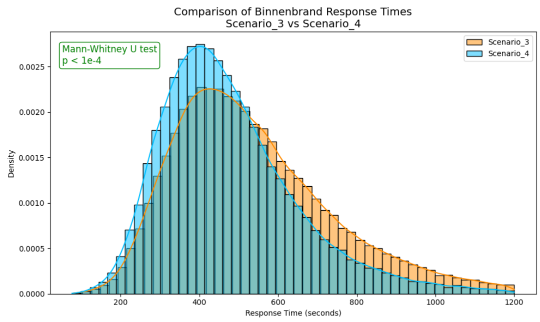

(b) Scenario 3 vs Scenario 4

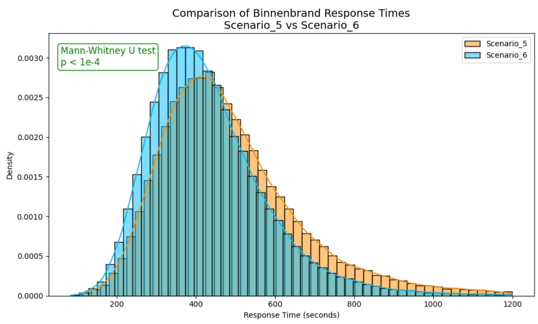

(c) Scenario 5 vs Scenario 6

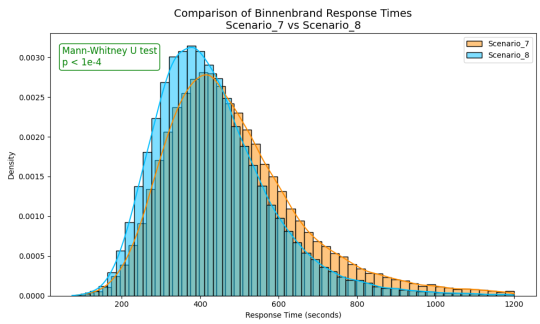

(d) Scenario 7 vs Scenario 8

Figure 7.12: Pairwise comparisons under scaled arrival for Indoor fire.

incidents occur very frequently without changes to the deployment pattern, saturation effects may arise. Such cases could delay treatment and increase fatality risk, although the proportions remain negligible (see Appendix A.4, Table A.4.0.6).

Compared with the normal-arrival experiments, Resuscitation FOTA in the old-deployment scenarios drops by about 11% under scaled arrival. By contrast, the new-deployment scenarios show stable FOTA with the new dispatch rule and only a minor decline with the old rule. The combination of the new deployment and the new dispatch rule therefore maintains performance even under high incident frequency.

Figures 7.13 and 7.14 show that high FOTA values are preserved in the new-deployment scenarios, even when incident frequency is doubled.

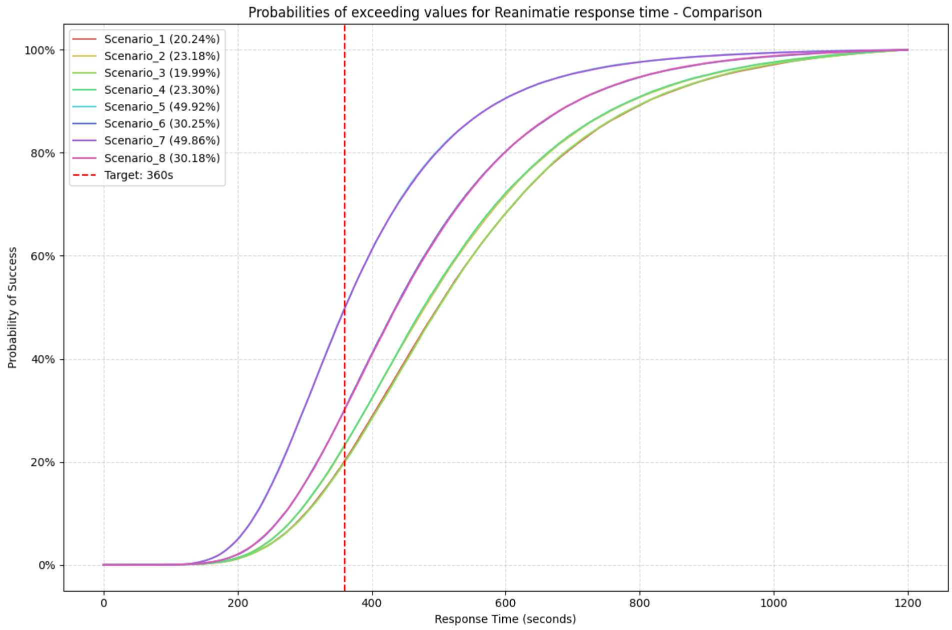

Figure 7.13: FOTA comparisons for Resuscitation under scaled arrival scenarios.

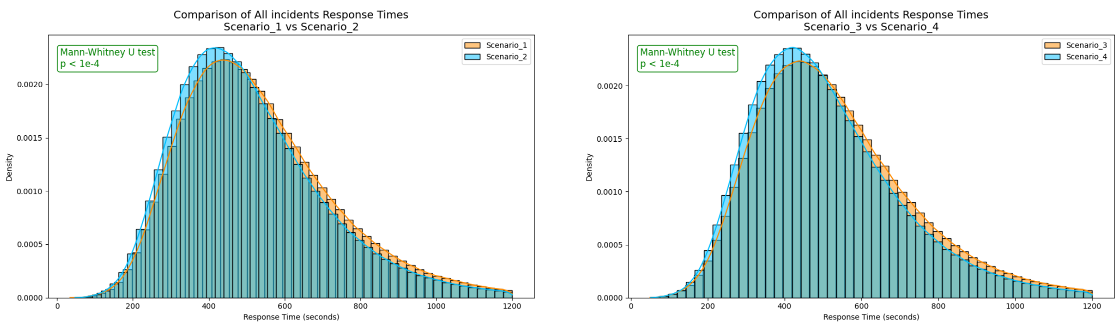

(a) Scenario 1 vs Scenario 2

(b) Scenario 3 vs Scenario 4

Figure 7.14: Pairwise comparisons of old-deployment scenarios under scaled arrival for Resuscitation.

7.2.3 Big-Incident Periods

We define *Big-incident periods* as intervals during which at least one major or very-major indoor fire is active. An interval ends once all large indoor fires within it are cleared.

7.2.3.1 Normal Arrival: Big-Incident Periods

All incidents (Big incident periods, normal arrival) External support use during Big-incident periods is rare, ranging from 0.00% to 0.02% across scenarios (Appendix A.4, Table A.4.0.7).

As shown in Figure 7.15, new-deployment scenarios (5–8) consistently achieve higher FOTA values than old-deployment scenarios (1–4), showing that the new layout maintains stronger coverage even during concentrated demand.

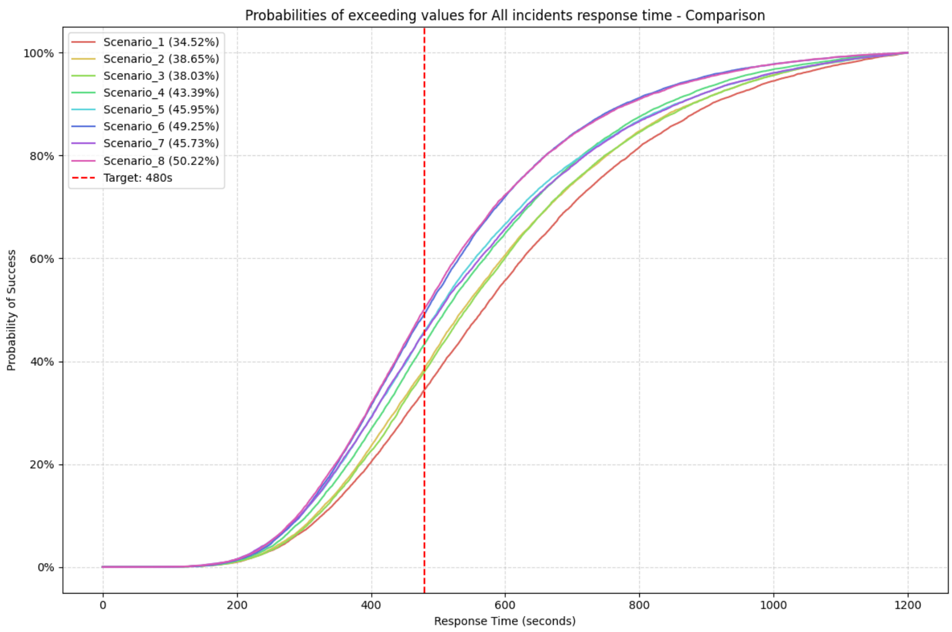

Figure 7.15: FOTA comparisons for all incidents during Big-incident periods under normal arrival rate.

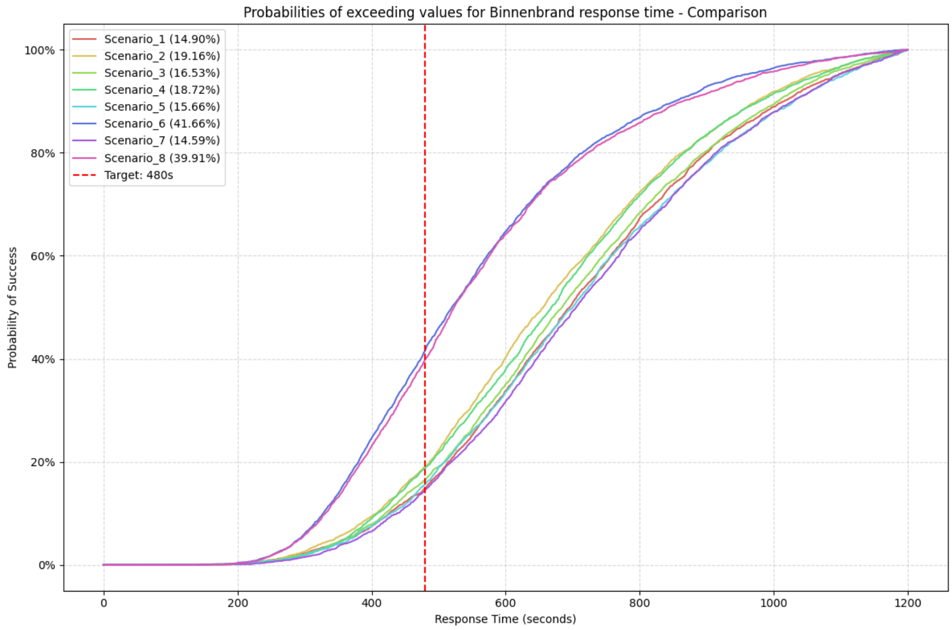

Figure 7.16: FOTA comparisons for Indoor fire during Big-incident periods under normal arrival rate.

Indoor fire (Big incident periods, normal arrival) For Indoor fire, external-support fractions remain low (0.00%–0.10%; see Appendix A.4, Table A.4.0.8).

Figure 7.16 shows FOTA comparisons for Indoor fire during Big-incident periods. Under the old deployment, relocation applied more often during Big-incident periods brings small gains (Scenario 1 vs Scenario 3), though these are not statistically significant, as illustrated in Figure 7.17. The new dispatch rule demonstrates clearer benefits in both deployment settings (Scenarios 2, 4, 6, and 8).

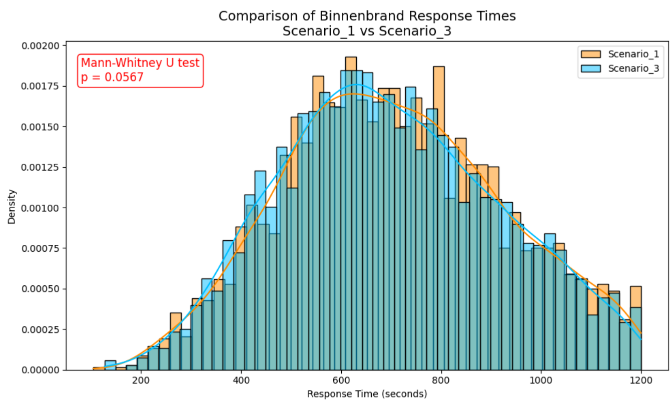

Figure 7.17: Scenario 1 vs Scenario 3 comparison for Indoor fire during Big-incident periods (normal arrival rate).

Resuscitation (Big incident periods, normal arrival) For Resuscitation, no external-support cases occur in any scenario during Big-incident periods (Appendix A.4, Table A.4.0.9).

Figure 7.18 shows FOTA comparisons for Resuscitation during Big-incident periods. Overall, new-deployment scenarios outperform old ones. Within the old deployment, the new dispatch rule generally improves results, whereas in the new deployment the old rule can sometimes yield higher FOTA values.

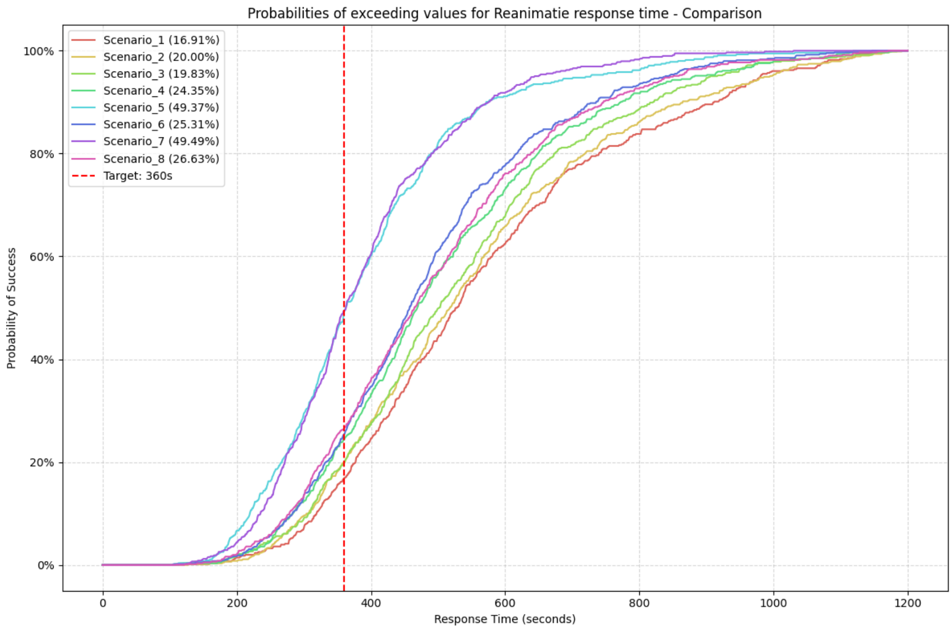

Figure 7.18: FOTA comparisons for Resuscitation during Big-incident periods under normal arrival rate.

7.2.3.2 Scaled Arrival: Big-Incident Periods

All incidents (Big incident periods, scaled arrival) External support during Big-incident periods under scaled arrivals remains uncommon, ranging from 0.02% to 0.28% (Appendix A.4, Table A.4.0.10).

Figure 7.19 presents FOTA comparisons for all incidents during Big-incident periods under scaled arrival scenarios. Compared with normal arrivals, old-deployment scenarios (1–4) show clear declines in FOTA (e.g. Scenario 2 drops about 7.5% relative to baseline). New-deployment scenarios show mixed patterns, with small increases in some cases and stability in others.

Interestingly, under scaled arrivals the share of Indoor fires during Big-incident periods falls from roughly one quarter to about one seventh of incidents. Since Indoor fires last

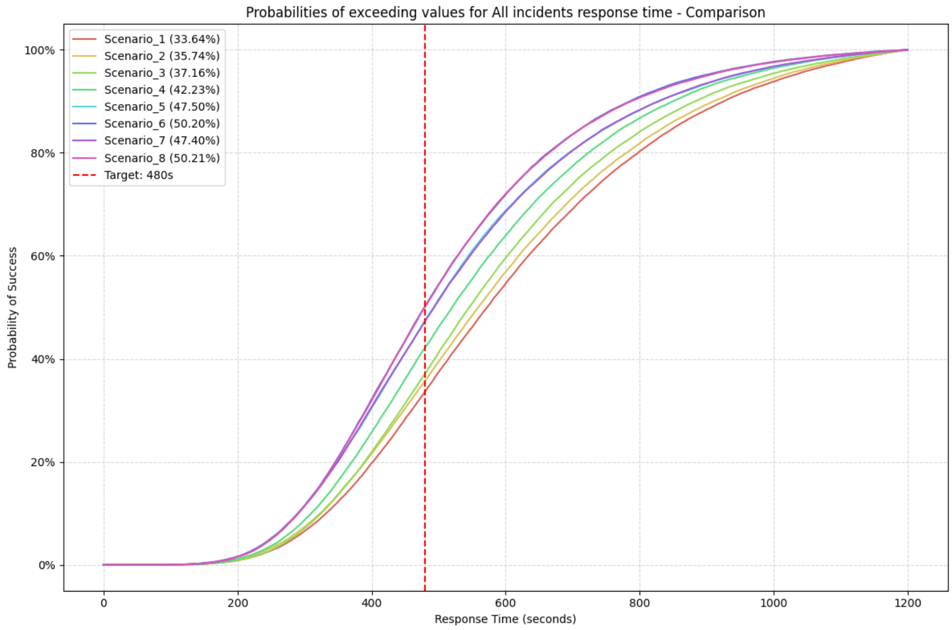

Figure 7.19: FOTA comparisons during Big-incident periods under scaled arrival rate. External-support occurrences increase overall, but remain low in absolute proportion.

longer and tie up vehicles, this shift frees more units for shorter calls, which can improve FOTA for those incidents. This effect appears consistently in our simulations, though the exact mechanism behind the change in incident composition under simple doubling of arrivals remains unclear.

Indoor fire (Big incident periods, scaled arrival) External-support fractions for Indoor fires during Big-incident periods under scaled arrivals range from 0.00% to 0.38% (Table A.4.0.11).

Figure 7.20 shows Indoor-fire FOTA under scaled arrival rates. Indoor-fire FOTA generally improves under scaled arrivals. In old-deployment scenarios, gains range from +1% to +21% across Scenarios 1–4; in new-deployment scenarios, some exceed +35% (Scenarios 5–8). These gains likely reflect the shift in incident-type composition together with effective relocations, especially in old-deployment Scenario 3, where relocation was executed in about 90% of triggers (Figure 7.21).

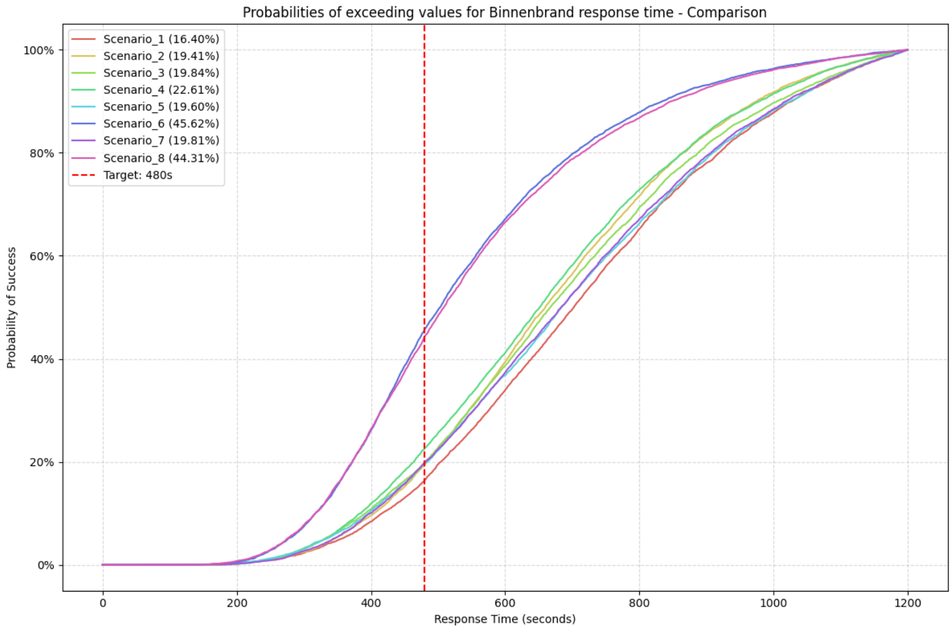

Figure 7.20: Indoor fire FOTA during Big-incident periods under scaled arrival rate. Many scenarios show marked increases compared to normal arrival rates.

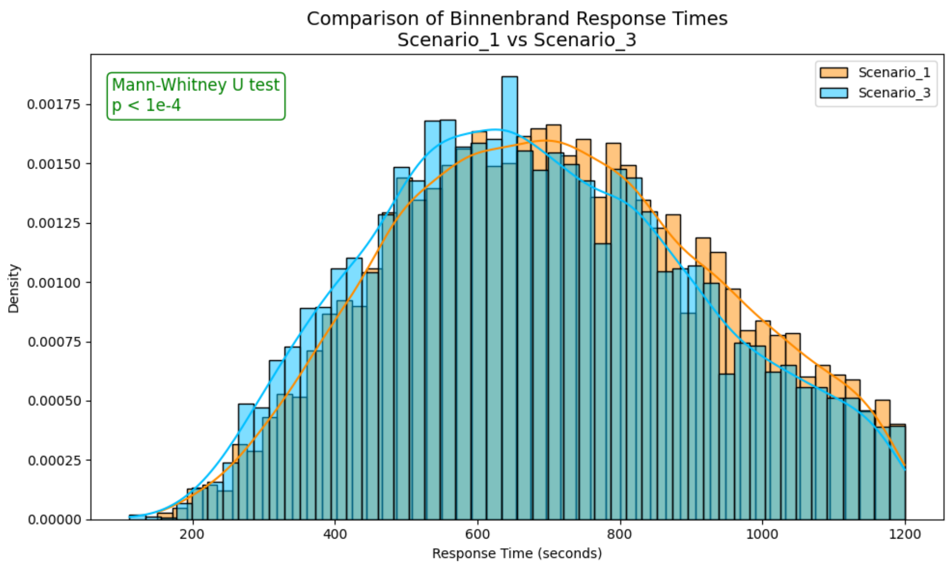

Figure 7.21: Comparison of response times between Scenario 1 and Scenario 3 during Big incident periods under scaled arrival. Scenario 3 shows significantly lower response times than Scenario 1.

Resuscitation (Big incident periods, scaled arrival) External-support fractions for Resuscitation remain very low (0.00%–0.22%; Appendix A.4, Table A.4.0.12). About half

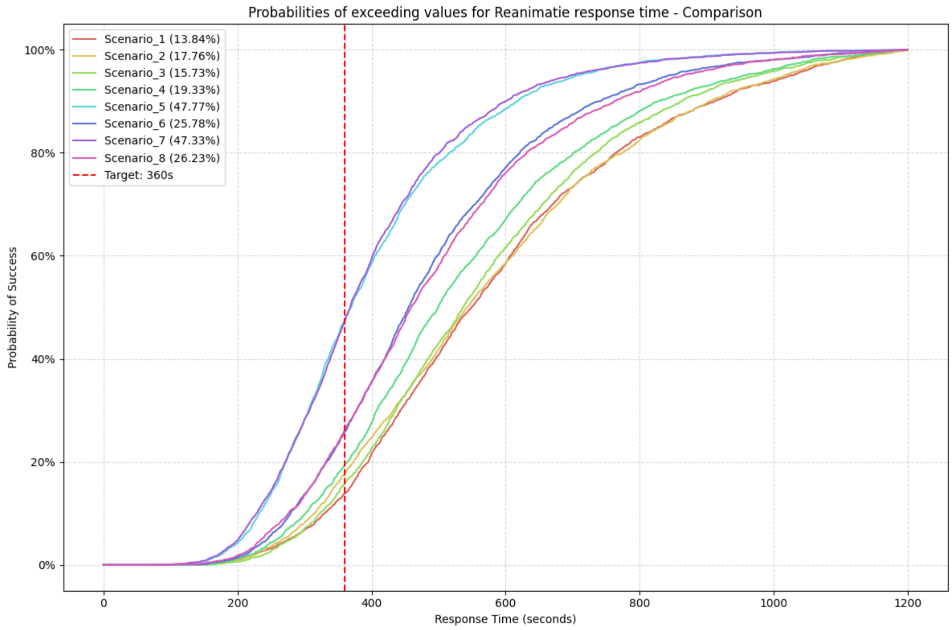

Figure 7.22: Resuscitation FOTA during Big incident periods under scaled arrival rate. New-deployment scenarios are generally more stable compared to old deployment.

of all such cases occur during Big-incident periods.

Figure 7.22 presents Resuscitation FOTA during Big-incident periods under scaled arrival rates. Compared with normal arrivals, Resuscitation FOTA during Big-incident periods under scaled arrivals declines by about 11%–20% in old-deployment scenarios (1–4). In new-deployment scenarios (5–8), the changes are smaller: –3.24%, +1.86%, –4.36%, and –1.50% respectively. The relative ranking between scenarios remains consistent with the normal-arrival case: new deployment generally outperforms old. For Resuscitation, the *new dispatch rule* gives better results under old deployment, while the *old dispatch rule* does better under the new deployment.

7.2.4 Overall Observations

This section summarizes the main insights derived from the simulation experiments across all incident types, deployment scenarios, and arrival conditions. The following points highlight key patterns and performance differences observed in the results:

- The new vehicle mix (Scenarios 5–8) generally outperforms the old deployment (Scenarios 1–4), even though the modeled crew numbers required to operate these vehicles are lower at several stations. This reflects more effective utilization of personnel.

- In most cases the *new dispatch rule* improves performance compared with the old rule, with the clearest gains observed for Indoor fire and All-incident FOTA.
- Resuscitation performance depends strongly on MOTO policy: reserving MOTOs at stations, rather than dispatching them broadly, yields better results in our simulations.
- The relocation algorithm requires refinement. Under the current simple trigger rule (activated by major/very-major indoor fires with immediate execution), frequent relocations can increase travel and return movements, worsening average response time. This points to the need for improved triggers (e.g., coverage thresholds, concurrency-based rules) and possibly a human-in-the-loop step for deciding relocations.

8

Conclusion

This research investigated how operational efficiency at the FDAA can be improved through more flexible deployment models. In particular, it examined the potential benefits of introducing variable unit compositions, including two- and four-person pumpers (TS2, TS4) and single-responder motorcycles (MOTO), and assessed the impact of adaptive dispatch rules and relocation strategies on response performance.

Given the impracticality and safety concerns of real-world experimentation, the study relied on an enhanced simulation framework. An existing simulator was extended to support flexible crew and vehicle configurations, composite units, and empirically grounded incident escalation. Adaptive dispatch rules and upgraded relocation algorithms—including adaptations of the MCRP and the LBAP—enabled more efficient allocation of resources under both routine and high-demand conditions. Additionally, the SDOM, a MILP-based approach, was introduced to optimize vehicle and crew allocation at the station level, considering personnel availability, vehicle capacity, and deployment constraints.

Simulation experiments demonstrated that adopting the new deployment model significantly improves performance. In this study, the main performance metric is the **Fraction of On-Time Arrivals (FOTA)**, which captures both response speed and effective coverage. FOTA measures the proportion of incidents reached within legally defined response-time thresholds; a higher FOTA indicates faster intervention and, indirectly, that incidents occur within adequately covered areas.

Without changing dispatch rules or using relocation algorithms, the introduction of smaller, faster units increased overall FOTA by 41.2% across all incidents. Applying the new dispatch rule alone, while keeping the current deployment, yielded an 11–13% improvement. For Indoor fire incidents, the new deployment outperformed the old, and the new dispatch rule further increased FOTA: 28–30% under the existing deployment and

15–17% under the new deployment. Notably, for Resuscitation incidents, the new deployment alone raised FOTA from 22.8% to 51.0%, a 124% relative improvement. The new dispatch rule improved performance by about 15% under the old deployment, but in the new deployment, careful MOTO allocation is necessary to avoid delays.

These improvements were consistent across high-demand scenarios and during major incidents. Even when incident frequency doubled, the new deployment maintained FOTA levels exceeding those of the current deployment under normal conditions, highlighting its resilience and efficiency.

Overall, the study shows that flexible deployment, combined with adaptive dispatch rules can substantially enhance emergency response performance. Optimized relocation strategies helped maintain geographic coverage, though their effectiveness depends on well-calibrated triggers to avoid unnecessary movements. Implementing these changes in practice would require adjustments to vehicle fleet composition, station-level allocation policies, and dispatch procedures. Future research could explore further tuning of relocation triggers, integration of dynamic traffic data, and real-time decision support to maximize the benefits of flexible deployment in operational settings.

9

Discussion and Recommendations

9.1 Discussion

The simulation framework developed in this study enabled systematic testing of new deployment models, dispatch rules, and relocation strategies. Several assumptions, however, limit its realism. Incident durations were adjusted mechanically once response times exceeded thresholds, which only approximates the consequences of late arrivals. Response time was also defined as the arrival of the first unit, even if this was a partially staffed vehicle. Such units may perform preparatory tasks, but their impact remains limited until larger crews arrive. A more refined model would assess the contribution of early arrivals in stabilizing incidents and their interaction with follow-up teams.

Travel speeds and relocation dynamics also require improvement. Current travel time factors are based on expert judgment rather than measured data. For TS-series trucks and buses, travel times are simple ratios of OSRM estimates, while drones are assumed much faster; for drones, future work could compute travel times based on straight-line distances and actual speed, as they are not constrained by roads. Relocation is determined solely by coverage gaps and routing distances. Differences in turnout times between professional and volunteer stations are ignored, which likely exaggerates relocation effectiveness in volunteer-based areas. Similarly, the dispatch algorithm evaluates stations in fixed order and treats all crews alike, overlooking turnout delays that can affect real performance. Incorporating differentiated turnout times into both relocation and dispatch would bring the model closer to operational practice.

The relocation algorithm is based on an optimization framework designed to improve coverage and resource allocation. Its practical implementation, however, requires careful consideration of operational factors beyond the model. For instance, staff scheduling, station familiarity, and training can influence the feasibility and effectiveness of dynamic relocation. Future work could explore integration with human-centered planning, ensuring that relocation strategies are compatible with routine operations and personnel readiness.

Challenges also arise when the fleet includes more diverse vehicle types. Introducing smaller units while reducing the number of high-capacity trucks makes it harder for relocation algorithms to maintain balanced coverage. Although the functionality is implemented, performance is constrained by vehicle availability and composition. Adaptive strategies that explicitly account for fleet mix and dynamic demand could better address this issue.

It is important to note that this study does not focus on reducing staff numbers. The fact that the new deployment configuration achieves better performance with fewer required crew does not imply downsizing, but rather a more effective use of available personnel. In practice, the department employs several times more firefighters than those directly matched to vehicles, and the results here should be seen as guidance for designing more efficient deployment schemes within existing resource constraints.

Finally, incident generation and escalation in the simulator depend on a limited dataset, risking underestimation of rare or emerging events. Continued data collection on escalation timings and spatial patterns would allow more robust modeling. A particular concern is cardiac arrest response. The present framework assumes that the immediate arrival of an AED-equipped motorcycle ensures survival, which is unrealistic. In practice, outcomes often rely on a two-stage sequence: rapid AED arrival followed by a crew capable of sustained resuscitation. Capturing this sequence is essential to evaluate the true role of motorcycle deployment.

9.2 Recommendation

Despite these limitations, the results show that flexible deployment models and adaptive dispatch rules can deliver meaningful improvements in response performance. For practice, it is recommended that the fire department continues to explore alternative fleet compositions, in particular the targeted use of AED-equipped motorcycles for cardiac arrests, where the performance gains are most evident. More flexible dispatch rules should also be adopted, as they can enhance efficiency with little additional cost.

Equally important is systematic data collection. Detailed records of escalation timings, turnout delays by station type, and resource utilization are needed to refine simulation models and support evidence-based decision making. By combining innovative deployment strategies with stronger data foundations, the fire department can further strengthen both the efficiency and resilience of its emergency response system, thereby enhancing safety and benefiting the public.

References

[1] Jordan L. Singer and Vincent N. Mosesso. **After the lights and sirens: Patient access delay in cardiac arrest**. *Resuscitation*, **155**:234–235, October 2020. 2, 14

[2] Pieter L. Van den Berg, Guido A. G. Legemaate, and Rob D. Van der Mei. **Increasing the Responsiveness of Firefighter Services by Relocating Base Stations in Amsterdam**. *Interfaces*, **47**(4):352–361, September 2017. 2, 21

[3] Dmitrii Usanov, Guido A. G. Legemaate, Peter M. Van de Ven, and Rob D. Van der Mei. **Fire truck relocation during major incidents**. *Naval Research Logistics (NRL)*, **66**(2):105–122, March 2019. 2, 20, 44

[4] Guido A. G. Legemaate, Jeffrey De Deijn, Sandjai Bhulai, and Rob D. Van der Mei. **Severe Weather-Based Fire Department Incident Forecasting**. *International Journal on Advances in Intelligent Systems*, **14**:14–23, 2021. 2, 21

[5] Joep Van den Bogaert. *The data-driven fire department.* Engd thesis, August 2019. 12, 14, 18, 25, 30, 31, 32, 35, 40, 54

[6] BBC News. **Motorbike pilot scheme for Merseyside firefighters**, 2010. Accessed: 2025-08-18. 15

[7] Singapore Civil Defence Force. **Emergency Medical Services**, 2025. Accessed: 2025-08-18. 15

[8] Lafayette County Fire Department. **Lafayette County Fire Department Adds New Pierce Mini-Pumper to Fleet**, 2025. Accessed: 2025-08-18. 15

[9] Martin Van Buuren, Rob D. Van der Mei, Karen Aardal, and Henk Post. **Evaluating dynamic dispatch strategies for emergency medical services: TIFAR simulation tool**. In *Proceedings of the 2012 Winter Simulation Conference (WSC)*, pages 1–12, Berlin, Germany, December 2012. IEEE. 17

[10] Constantine Toregas, Ralph Swain, Charles ReVelle, and Lawrence Bergman. **The Location of Emergency Service Facilities**. *Operations Research*, **19**(6):1363–1373, October 1971. 19

[11] Richard Church and Charles ReVelle. **The maximal covering location problem**. *Papers of the Regional Science Association*, **32**(1):101–118, December 1974. 19

[12] Peter Kolesar and Warren E. Walker. **An Algorithm for the Dynamic Relocation of Fire Companies**. *Operations Research*, **22**(2):249–274, April 1974. 19

[13] Bohuš Leitner, Michal Ballay, and Marek Kvet. **Optimization of Fire Brigade Deployment by Means of Mathematical Programming**. *Processes*, **11**(4):1262, April 2023. 22

[14] Juan Pérez, Sebastián Maldonado, and Héctor López-Ospina. **A fleet management model for the Santiago Fire Department**. *Fire Safety Journal*, **80**:62–71, 2016. 23

[15] Sean J. Taylor and Benjamin Letham. **Forecasting at scale**, September 2017. 32

[16] Travis E. Oliphant. **A Bayesian perspective on estimating mean, variance, and standard-deviation from data**. *All Faculty Publications*, (278):1–17, December 2006. 37

[17] Remy Stieglis, Bas J. Verkaik, Hanno L. Tan, Rudolph W. Koster, Hans Van Schuppen, and Christian Van der Werf. **Association Between Delay to First Shock and Successful First-Shock Ventricular Fibrillation Termination in Patients With Witnessed Out-of-Hospital Cardiac Arrest**. *Circulation*, **151**(3):235–244, January 2025. 67

Appendix A

Appendix: Full Tables

A.1 Dispatch Rule Table

Incident Type	RV	HV	WO	TS2	TS4	TS6	TS6_P	MOTO	DRONE
General assistance		1			1				
Police assistance					1				
Alarm	1						1		
Ambulance assistance					1				
Fire alarm	1						1		
Animal assistance			1		1				
Outdoor fire					1				
Resuscitation									
Elevator entrapment									
Person in water			1		1				
Hazardous substances				1					
Indoor fire							1		
Extreme weather	1				1				
Other fire					1				
Vehicle in water			1		1				
Water rescue			1		1				
Animal in water			1		1				
Aviation rescue						1			
Unknown				1					

Table A.1: Vehicle dispatch table for Medium incident types

Incident Type	**RV**	**HV**	**WO**	**TS2**	**TS4**	**TS6**	**TS6_P**	**MOTO**	**DRONE**
General assistance		1			1				
Police assistance					1				
Alarm	1					1			
Ambulance assistance					1				
Fire alarm	1					1			
Animal assistance			1		1				
Outdoor fire					1				
Resuscitation									
Elevator entrapment									
Person in water			1		1				
Hazardous substances				1					
Indoor fire	1					1			
Extreme weather	1				1				
Other fire					1				
Vehicle in water		1	1		1				
Water rescue			1		1				
Animal in water			1		1				
Aviation rescue		1				1			
Unknown				1					

Table A.2: Vehicle dispatch table for Major incident types

Incident Type	RV	HV	WO	TS2	TS4	TS6	TS6_P	MOTO	DRONE
General assistance					1				
Police assistance					1				
Alarm						1			
Ambulance assistance					1				
Fire alarm						1			
Animal assistance			1		1				
Outdoor fire					1				
Resuscitation									
Elevator entrapment									
Person in water			1		1				
Hazardous substances				1					
Indoor fire						1			
Extreme weather	1				1				
Other fire					1				
Vehicle in water			1		1				
Water rescue			1		1				
Animal in water			1		1				
Aviation rescue						1			
Unknown				1					

Table A.3: Vehicle dispatch table for Very Major incident types

A.2 Legacy Deployment Configuration

Table A.4: Legacy deployment configuration: Only TS6 vehicles, with current staffing patterns.

Station	BUS	MOTO	TS2	TS4	TS6	TS_crew_ft	TS_crew_pt
A	1	0	0	0	2	6	6
B	1	0	0	0	1	6	0
C	1	0	0	0	1	6	0
D	1	0	0	0	1	6	0
E	1	0	0	0	1	6	0
F	1	0	0	0	1	6	0
G	1	0	0	0	2	0	12
H	1	0	0	0	1	6	0
I	1	0	0	0	1	6	0
J	1	0	0	0	1	0	6
K	1	0	0	0	1	6	0
L	1	0	0	0	1	6	0
M	1	0	0	0	1	6	0
N	1	0	0	0	1	6	0
O	1	0	0	0	2	0	12
P	1	0	0	0	2	6	6
Q	1	0	0	0	1	0	6
R	1	0	0	0	1	6	0

A.3 New Deployment Configuration

Table A.5: New deployment configuration: Incorporating MOTO, TS2, and TS4 vehicles with reduced staffing.

Station	BUS	MOTO	TS2	TS4	TS6	TS_crew_ft	TS_crew_pt
A	1	1	1	1	2	8	6
B	1	1	1	1	0	6	0
C	1	1	0	1	0	4	0
D	1	1	0	1	1	7	0
E	1	1	1	1	1	8	0
F	1	1	1	0	1	8	0
G	1	1	1	1	0	0	6
H	1	1	1	0	1	8	0
I	1	1	0	1	0	4	0
J	1	1	1	0	0	0	3
K	1	1	1	0	1	6	0
L	1	1	0	1	0	4	0
M	1	1	0	1	0	4	0
N	1	1	1	0	0	3	0
O	1	1	0	1	1	0	6
P	1	1	1	0	0	3	3
Q	1	1	0	1	0	0	4
R	1	1	0	1	1	7	0

A.4 External-Support Fractions (Full Listing)

Below we reproduce the **exact** external-support fraction rows as they were computed in the experiments. These are grouped by experiment (normal / scaled arrival rates) and by period (All periods / Big-incident periods) and by incident type (All incidents / Indoor fire / Resuscitation). Use these tables for auditing and reproducibility.

A.4.0.1 Normal Arrival Rate — All Periods

```
All incidents (All periods, normal arrival rate):
Scenario_1: All Incidents EXTERNAL fraction = 0 / 864172 = 0.00%
Scenario_2: All Incidents EXTERNAL fraction = 2 / 853961 = 0.00%
Scenario_3: All Incidents EXTERNAL fraction = 1 / 865343 = 0.00%
Scenario_4: All Incidents EXTERNAL fraction = 1 / 853606 = 0.00%
Scenario_5: All Incidents EXTERNAL fraction = 1 / 852541 = 0.00%
Scenario_6: All Incidents EXTERNAL fraction = 2 / 858787 = 0.00%
```

```
Scenario_7: All Incidents EXTERNAL fraction = 3 / 855151 = 0.00%
Scenario_8: All Incidents EXTERNAL fraction = 2 / 857282 = 0.00%
```

A.4.0.2 Normal Arrival Rate — Indoor Fire

```
Scenario_1: Indoor Fire EXTERNAL fraction = 0 / 43612   = 0.00%
Scenario_2: Indoor Fire EXTERNAL fraction = 0 / 33204   = 0.00%
Scenario_3: Indoor Fire EXTERNAL fraction = 1 / 43394   = 0.00%
Scenario_4: Indoor Fire EXTERNAL fraction = 0 / 33063   = 0.00%
Scenario_5: Indoor Fire EXTERNAL fraction = 1 / 33147   = 0.00%
Scenario_6: Indoor Fire EXTERNAL fraction = 2 / 36896   = 0.01%
Scenario_7: Indoor Fire EXTERNAL fraction = 3 / 32940   = 0.01%
Scenario_8: Indoor Fire EXTERNAL fraction = 1 / 36826   = 0.00%
```

A.4.0.3 Normal Arrival Rate — Resuscitation

```
Scenario_1: Resuscitation EXTERNAL fraction = 0 / 45804 = 0.00%
Scenario_2: Resuscitation EXTERNAL fraction = 0 / 45866 = 0.00%
Scenario_3: Resuscitation EXTERNAL fraction = 0 / 45706 = 0.00%
Scenario_4: Resuscitation EXTERNAL fraction = 0 / 46146 = 0.00%
Scenario_5: Resuscitation EXTERNAL fraction = 0 / 45614 = 0.00%
Scenario_6: Resuscitation EXTERNAL fraction = 0 / 46017 = 0.00%
Scenario_7: Resuscitation EXTERNAL fraction = 0 / 46389 = 0.00%
Scenario_8: Resuscitation EXTERNAL fraction = 0 / 46390 = 0.00%
```

A.4.0.4 Scaled Arrival Rate — All Periods

```
Scenario_1: All Incidents EXTERNAL fraction = 26 / 1721681 = 0.00%
Scenario_2: All Incidents EXTERNAL fraction = 242 / 1700589 = 0.01%
Scenario_3: All Incidents EXTERNAL fraction = 53 / 1720425 = 0.00%
Scenario_4: All Incidents EXTERNAL fraction = 149 / 1700058 = 0.01%
Scenario_5: All Incidents EXTERNAL fraction = 53 / 1707184 = 0.00%
Scenario_6: All Incidents EXTERNAL fraction = 96 / 1716668 = 0.01%
Scenario_7: All Incidents EXTERNAL fraction = 55 / 1705651 = 0.00%
Scenario_8: All Incidents EXTERNAL fraction = 73 / 1717256 = 0.00%
```

A.4.0.5 Scaled Arrival Rate — Indoor Fire

```
Scenario_1: Indoor Fire EXTERNAL fraction = 3 / 86724  = 0.00%
Scenario_2: Indoor Fire EXTERNAL fraction = 28 / 66664 = 0.04%
Scenario_3: Indoor Fire EXTERNAL fraction = 13 / 86512 = 0.02%
Scenario_4: Indoor Fire EXTERNAL fraction = 32 / 66571 = 0.05%
Scenario_5: Indoor Fire EXTERNAL fraction = 24 / 65790 = 0.04%
Scenario_6: Indoor Fire EXTERNAL fraction = 36 / 75015 = 0.05%
Scenario_7: Indoor Fire EXTERNAL fraction = 32 / 66313 = 0.05%
Scenario_8: Indoor Fire EXTERNAL fraction = 32 / 75164 = 0.04%
```

A.4.0.6 Scaled Arrival Rate — Resuscitation

```
Scenario_1: Resuscitation EXTERNAL fraction = 0 / 91420 = 0.00%
Scenario_2: Resuscitation EXTERNAL fraction = 11 / 91224 = 0.01%
Scenario_3: Resuscitation EXTERNAL fraction = 3 / 90756 = 0.00%
Scenario_4: Resuscitation EXTERNAL fraction = 10 / 91860 = 0.01%
Scenario_5: Resuscitation EXTERNAL fraction = 0 / 92257 = 0.00%
Scenario_6: Resuscitation EXTERNAL fraction = 0 / 91717 = 0.00%
Scenario_7: Resuscitation EXTERNAL fraction = 0 / 92711 = 0.00%
Scenario_8: Resuscitation EXTERNAL fraction = 0 / 91706 = 0.00%
```

A.4.0.7 Big-Incident Periods — All Incidents (Normal Arrival)

```
Scenario_1: 0 / 13123  = 0.00%
Scenario_2: 1 / 14407  = 0.01%
Scenario_3: 1 / 13636  = 0.01%
Scenario_4: 1 / 14692  = 0.01%
Scenario_5: 1 / 13136  = 0.01%
Scenario_6: 2 / 13931  = 0.01%
Scenario_7: 3 / 13007  = 0.02%
Scenario_8: 1 / 13922  = 0.01%
```

A.4.0.8 Big-Incident Periods — Indoor Fire (Normal Arrival)

```
Scenario_1: 0 / 3101  = 0.00%
Scenario_2: 0 / 3148  = 0.00%
Scenario_3: 1 / 3158  = 0.03%
```

```
Scenario_4: 0 / 3135  = 0.00%
Scenario_5: 1 / 2856  = 0.04%
Scenario_6: 2 / 3282  = 0.06%
Scenario_7: 3 / 2888  = 0.10%
Scenario_8: 1 / 3182  = 0.03%
```

A.4.0.9 Big-Incident Periods — Resuscitation (Normal Arrival)

```
Scenario_1: 0 / 556  = 0.00%
Scenario_2: 0 / 630  = 0.00%
Scenario_3: 0 / 605  = 0.00%
Scenario_4: 0 / 616  = 0.00%
Scenario_5: 0 / 553  = 0.00%
Scenario_6: 0 / 648  = 0.00%
Scenario_7: 0 / 586  = 0.00%
Scenario_8: 0 / 612  = 0.00%
```

A.4.0.10 Big-Incident Periods — All Incidents (Scaled Arrival)

```
Scenario_1: 10 / 46277  = 0.02%
Scenario_2: 139 / 49023 = 0.28%
Scenario_3: 32 / 50121  = 0.06%
Scenario_4: 79 / 52847  = 0.15%
Scenario_5: 26 / 50099  = 0.05%
Scenario_6: 49 / 53818  = 0.09%
Scenario_7: 30 / 51732  = 0.06%
Scenario_8: 45 / 53976  = 0.08%
```

A.4.0.11 Big-Incident Periods — Indoor Fire (Scaled Arrival)

```
Scenario_1: 0 / 7116  = 0.00%
Scenario_2: 17 / 6985 = 0.24%
Scenario_3: 10 / 7347 = 0.14%
Scenario_4: 27 / 7190 = 0.38%
Scenario_5: 18 / 6459 = 0.28%
Scenario_6: 21 / 7534 = 0.28%
Scenario_7: 22 / 6785 = 0.32%
Scenario_8: 25 / 7596 = 0.33%
```

A.4.0.12 Big-Incident Periods — Resuscitation (Scaled Arrival)

```
Scenario_1: 0 / 2225 = 0.00%
Scenario_2: 5 / 2313  = 0.22%
Scenario_3: 1 / 2391  = 0.04%
Scenario_4: 5 / 2571  = 0.19%
Scenario_5: 0 / 2332  = 0.00%
Scenario_6: 0 / 2634  = 0.00%
Scenario_7: 0 / 2544  = 0.00%
Scenario_8: 0 / 2604  = 0.00%
```