

SIMULATION-BASED EVALUATION OF SIGNALISED INTERSECTIONS FOR EMERGENCY SAFETY LANDINGS IN URBAN AIR MOBILITY

Pavithra Sathya Kumar[1], Michael Schultz[1] & Ehsan Asadi[1]

[1]Chair of Air Traffic Concepts , University of the Bundeswehr Munich , Munich, Germany

Abstract

Emergency safety landings of Advanced air mobility (AAM) vehicles in dense urban environments may require rapid coordination between aerial contingency management and ground traffic control. While current AAM concepts primarily rely on designated vertiports and predefined alternates, local contingencies may require the assessment of non-designated urban infrastructure. This paper presents an initial proof-of-concept evaluation of the ground-side traffic-management feasibility of using signalised urban intersections as temporary emergency safety landing-site candidates within an AGTM framework.

A microscopic traffic simulation framework is implemented in SUMO and controlled through TraCI for a peripheral urban road network in Munich. Five scenarios are compared: baseline operation (S0), uncoordinated all-red emergency closure (S1), fixed-site coordinated rerouting (S2), traffic-aware dynamic site selection (S3), and AAM-constrained dynamic selection (S4). Candidate intersections are ranked using live traffic indicators, including local vehicle counts, halted vehicles, route conflicts, and simplified ground-side controllability proxies based on controlled-lane coverage and signal-state length.

The results show that temporary intersection protection is not a simple traffic-light override problem. Fixed-site rerouting provides only partial mitigation and can transfer congestion to alternative corridors. Dynamic selection improves performance, but reveals a trade-off between local landing-site clearance and network-level resilience. Traffic-aware dynamic selection achieves the strongest local clearance by favouring smaller, lower-disruption intersections; in contrast, AAM-constrained selection accepts a higher local queue but selects larger intersections according to controllability proxies. This improves network-level mitigation, including reduced halted vehicles during the protection interval and lower recovery-phase waiting-time values. The findings show that AGTM can support the assessment of signalised intersections as controllable ground-side candidates for AAM emergency landing preparation, while physical landing feasibility remains a separate certification problem.

Keywords: Advanced Air Mobility ; Emergency Landing; Advanced Ground Traffic Management; SUMO; TraCI; Dynamic Rerouting

1. Introduction

Advanced Air Mobility introduces contingency-management requirements that differ from those of conventional ground or air transport. Emergency landing decisions for electric vertical take-off and landing (eVTOL) vehicles cannot be evaluated only by airborne reachability. In dense urban environments, the selected ground location must also be accessible, controllable, and recoverable under active traffic conditions.

In conventional aviation, diversion planning is formalised through Extended Diversion Time Operations (EDTO [1]), which require aircraft to remain within reach of suitable alternate aerodromes [2]. This principle cannot be transferred directly to AAM. Urban landing infrastructure is spatially constrained, designated vertiports may be occupied or unavailable, and non-designated locations may be obstructed by vehicles, pedestrians, street furniture, or local access restrictions. Current AAM and U-space concepts recognise contingency management and predefined alternates, but they do not yet provide an operational mechanism for preparing a non-designated urban ground location during a time-critical safety landing request [3, 4].

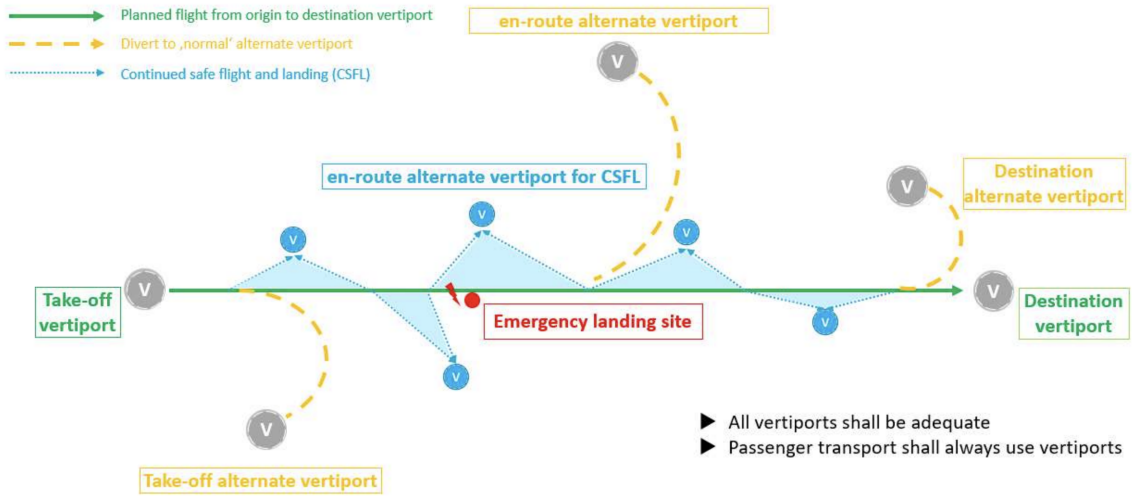


Figure 1 – Alternate vertiports distributed along planned AAM trajectories within standard aviation routing frameworks [4].

This paper examines signalised urban intersections as temporary ground-side candidates for emergency AAM safety landing preparation. The study does not claim that intersections are aeronautically certified landing sites. It addresses a narrower traffic-management question: if a signalised intersection is selected as a candidate ground area, can it be temporarily cleared, restricted, and managed while limiting disruption to the surrounding road network? This distinction is important because physical landing feasibility and operational traffic feasibility are separate problems. Aircraft-side trajectory planning and landing control determine whether an eVTOL can reach and execute a landing manoeuvre, whereas ground-side traffic management determines whether the selected area can be made available and protected.

Signalised intersections are relevant to this question because they are controllable road-network nodes. Their signal states can be modified, approaching vehicles can be monitored, and traffic can be redirected before a protection interval begins. Previous work has proposed the road network as a potential safety net for urban air traffic, where intersections may serve as emergency landing areas if supported by traffic-management measures [5]. Building on this concept, the present study evaluates the traffic consequences of temporarily converting selected signalised intersections into protected candidate areas.

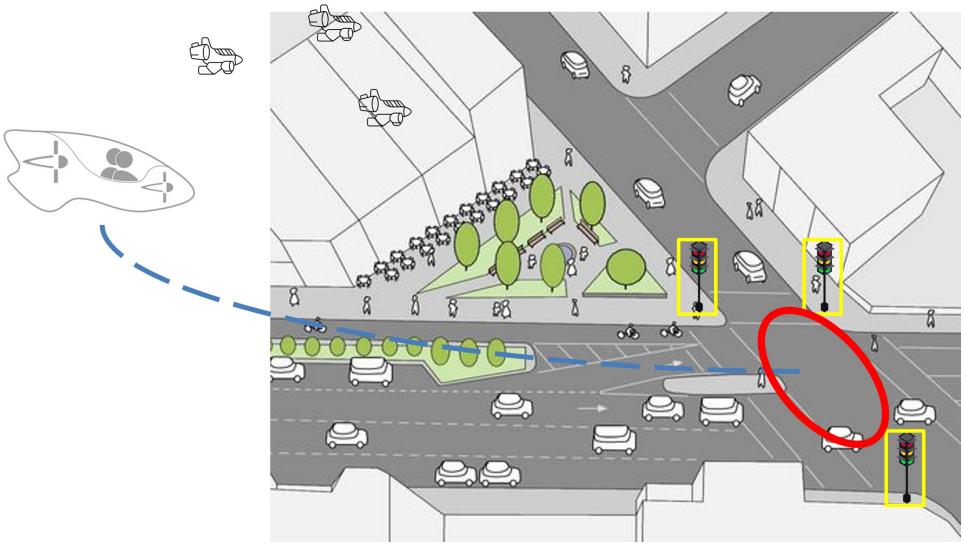


Figure 2 – Conceptual integration of a signalised urban intersection as a temporary safety landing-site candidate managed under active ground traffic control.

The operational difficulty is that protecting an intersection imposes an abrupt capacity reduction at a network node. A locally available candidate may therefore generate queue spillback, delay transfer to neighbouring corridors, and delayed recovery after the restriction is released. For this reason, emergency landing-site preparation must be treated as an integrated air–ground traffic-management

problem. The ground-side decision is not only which location can be used, but also how the surrounding traffic system responds to its temporary restriction.

This paper addresses the following research question: to what extent can Advanced Ground Traffic Management (AGTM) support the temporary restriction and preparation of signalised intersections for emergency AAM safety landing requests while limiting local and network-level disruption?

An initial version of this concept was introduced in earlier work on multimodal traffic coordination for safety landings [6]. The present study extends that initial concept by developing and evaluating a microscopic SUMO/TraCI simulation framework for AGTM-supported intersection restriction, rerouting, dynamic candidate selection, and recovery analysis. A microscopic SUMO/TraCI simulation framework is developed for an urban road network in Munich. The AAM emergency is represented as a time-critical AGTM event that triggers temporary restriction of a selected signalised intersection. Five scenarios are compared: baseline operation, uncoordinated all-red closure, fixed-site coordinated rerouting, traffic-aware dynamic site selection, and AAM-constrained dynamic selection. The scenarios are evaluated using network-level indicators, including average speed, halted vehicles, accumulated waiting time, and recovery behaviour, together with local indicators near the selected intersection.

The AAM-constrained strategy uses controlled-lane coverage and signal-state length as simplified ground-side controllability proxies. These variables do not represent aircraft dynamics, touchdown geometry, obstacle clearance, downwash, or certification requirements. They are used only to distinguish small low-impact junctions from larger traffic-control nodes that may be more relevant for emergency landing-site preparation.

The contribution of this paper is threefold:

a. It develops a SUMO/TraCI-based framework for modelling temporary ground-side preparation of signalised intersections during emergency AAM safety landing requests.
b. It quantifies the local and network-level impact of uncoordinated all-red intersection closure and fixed-site coordinated rerouting.
c. It compares traffic-aware and AAM-constrained dynamic selection strategies to identify the trade-off between local landing-area clearance and network-level resilience.

The remainder of this paper is structured as follows. Section 2. reviews AAM contingency concepts, emergency landing-site studies, and signalised intersection control. Section 3. presents the simulation framework, AGTM control logic, and scenario design. Section 4. reports the results. Section 5. discusses operational implications and limitations. Section 6. concludes the paper.

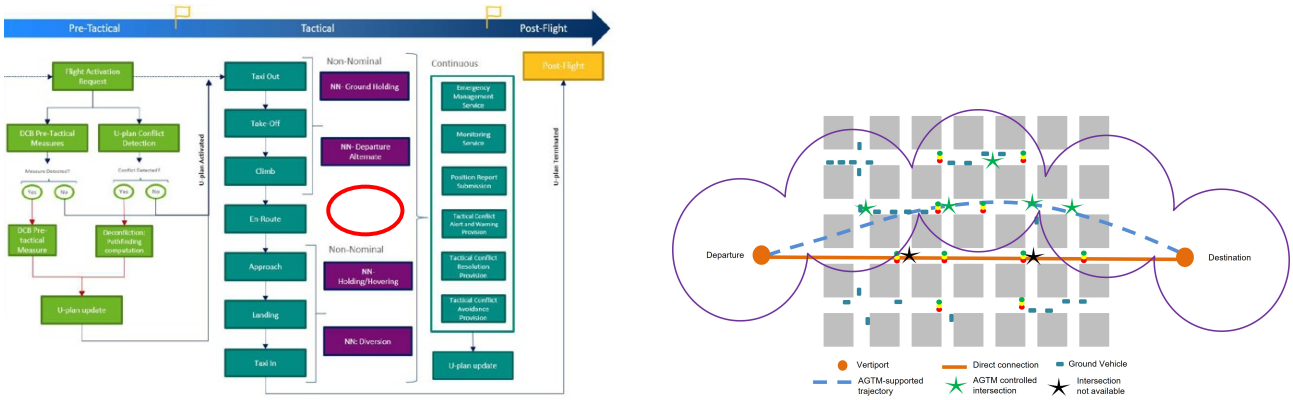
2. Related Work

2.1 AAM Contingency Management and Diversion Concepts

AAM traffic-management concepts provide the operational basis for planned eVTOL operations, but their treatment of emergency ground preparation remains limited. The FAA/NASA AAM Concept of Operations describes stakeholder roles, information exchange, operational coordination, and traffic-management services for increasingly complex AAM operations [3]. The SESAR U-space Concept of Operations addresses airspace structure, vertiport operations, flight rules, automation levels, and contingency handling for European AAM integration [4]. These concepts establish the need for structured contingency management, but they primarily focus on airspace coordination, vertiport operations, and predefined alternate options.

The limitation of the present problem is the treatment of the ground location. Existing concepts do not specify how a non-designated urban road location should be cleared, restricted, protected, and released during a time-critical safety landing. Intelligent contingency-management frameworks for AAM emphasise the need for autonomous or semi-autonomous decision-making under uncertainty, including vehicle-state assessment, environmental monitoring, and mission modification [7]. U-space contingency-management studies similarly identify technical failures, communication disruptions, environmental disturbances, and infrastructure unavailability as relevant contingency sources [8]. However, these frameworks do not quantify the road-network consequences of preparing a temporary landing area within active urban traffic.

Recent real-time landing-management work addresses the allocation of eVTOLs to landing pads under uncertain and high-density operations [9]. Such studies are relevant because they demonstrate the importance of real-time landing-resource assignment. Nevertheless, landing-pad assignment



(a) Decision layers in SESAR U-space ConOps, highlighting the non-designated infrastructure gap [4].

(b) Trajectory contingency management interfacing with AGTM-managed intersection candidates.

Figure 3 – Airspace contingency management context, contrasting current U-space decision-layer boundaries with an integrated air–ground landing-site preparation approach.

differs from the problem studied here: designated pads are assumed to be part of the landing infrastructure, whereas signalised intersections must first be made temporarily available through ground traffic control. The unresolved interface is therefore the transition from airborne contingency decision-making to operational preparation of a road-network location.

2.2 Emergency Landing and Landing-Site Feasibility

Emergency landing research for AAM can be grouped into three main streams: trajectory planning, aircraft-side landing control, and landing-site identification. Emergency landing trajectory planning has been formulated as an assignment and optimisation problem in which multiple eVTOL vehicles are routed to suitable landing options under response-time, priority, and separation constraints [10]. This work is central for airborne contingency planning, but it does not model the traffic-management process required to secure a selected ground location.

Aircraft-side control studies address the execution of emergency landings under degraded conditions. For example, emergency gliding guidance for tilt-wing eVTOL aircraft has been developed for power-loss scenarios in obstacle-constrained urban environments using control barrier functions [11]. Related autonomous self-landing methods for AAM address abnormal situations such as GNSS degradation, loss of control link, low battery, and obstacle avoidance [12]. These studies improve the safety of the descent and touchdown process. However, they do not determine whether the intended touchdown area can be cleared of vehicles or protected from new traffic access.

Landing-site identification studies add the physical and spatial dimension. Safe landing-zone detection surveys show that candidate selection may depend on terrain, obstacles, surface properties, perception accuracy, and environmental uncertainty [13]. Passenger-drone applications for emergency medical services use terrain and obstacle data to identify feasible landing locations while considering landing-area dimensions, safety areas, obstacle-free volumes, and approach/departure corridors [14]. Such criteria are necessary for physical feasibility assessment. However, they typically treat site availability as a spatial condition, whereas a signalised intersection is a dynamic traffic node whose availability changes with demand, queues, signal states, and rerouting behaviour.

The closest prior concept to the present study is the use of intermodal ground infrastructure as a safety net for AAM operations [5]. That work proposes that roads and intersections can support emergency landing options when combined with adaptable ground traffic management. It establishes the conceptual basis for treating roads and intersections as controllable emergency-support infrastructure. The present paper extends this line of work through a microscopic SUMO/TraCI evaluation of emergency intersection restriction, fixed-site rerouting, traffic-aware dynamic selection, and AAM-constrained candidate filtering in an urban Munich road network.

2.3 Signalised Intersection Control and AGTM

Signalised intersections provide an existing mechanism for controlling access to urban road-network nodes. Adaptive traffic signal control uses real-time traffic states to modify phase sequences, green

times, offsets, or coordination between neighbouring intersections. Multi-intersection adaptive control has been shown to reduce congestion, waiting time, and stops compared with fixed-time, actuated, or green-wave strategies [15]. Real-time safety-oriented signal control further shows that adaptive traffic signals can be optimised for safety-related objectives rather than only for delay or throughput [16].
Microscopic simulation is commonly used to evaluate such control strategies because it captures vehicle-level movement, queues, waiting time, and route-dependent effects. SUMO-based studies demonstrate how adaptive traffic-light logic can be implemented and evaluated using operational indicators such as lane density, vehicle counts, waiting time, and time loss [17]. AAM-SUMO further shows that SUMO can be extended to examine interactions between AAM concepts and ground transportation systems at a large scale [18]. In the present study, TraCI is used to retrieve traffic-light states, controlled lanes, vehicle routes, vehicle speeds, and waiting-time indicators, and to implement runtime signal restrictions and rerouting interventions [19, 20].
Emergency vehicle priority and preemption studies provide an important comparison case. Connected-vehicle-based signal priority methods aim to reduce emergency vehicle response time by providing preempted green phases and managing communication delays [21]. Traffic signal preemption and prioritisation protocols similarly modify signal states to support priority vehicles in dense road networks [22]. These approaches show that signalised intersections can support safety-critical interventions. However, the AAM safety-landing case reverses the control objective. Emergency vehicle preemption keeps the intersection available as a passage node for a priority vehicle. Emergency AAM landing preparation requires the intersection to be temporarily removed from normal traffic use and converted into a protected area.
The reviewed literature provides methods for airborne contingency management, emergency landing trajectory planning, aircraft-side landing control, spatial landing-site identification, and adaptive signal control. However, these strands do not quantify the road-network response when a signalised intersection is temporarily converted into a protected AAM safety-landing candidate. This paper addresses that gap by evaluating uncoordinated all-red closure, fixed-site pre-closure rerouting, traffic-aware dynamic site selection, and AAM-constrained candidate filtering within a SUMO/TraCI-based AGTM framework.

3. Methodology

3.1 Simulation Framework and Study Area

This study evaluates the ground-traffic feasibility of temporarily restricting signalised intersections as candidate safety landing areas for Advanced Air Mobility (AAM). The analysis is conducted in the microscopic traffic simulator SUMO and controlled at runtime through the Traffic Control Interface (TraCI). SUMO is used to model the road network, vehicle demand, vehicle movement, and traffic-light systems, while TraCI provides access to traffic-light states, controlled lanes, vehicle routes, speeds, halted vehicles, and waiting-time indicators.
The study area is an urban road network in Munich, Germany. It contains arterial corridors, local access links, and multiple signalised intersections with different geometric scales and traffic loads. This network structure enables comparison between fixed-site and dynamic intersection-selection strategies under identical traffic-demand and network-topology conditions. Across all scenarios, the network, demand file, vehicle composition, simulation horizon, and event timing are held constant. Therefore, differences between scenarios are attributable to the AGTM intervention logic rather than changes in traffic demand or network configuration.
The simulation horizon is $T = 3600$ s. The AGTM decision time is set to $t_{\mathrm{d}} = 900$ s, after the network has reached active traffic conditions. The closure begins at $t_{\mathrm{c}} = 1200$ s and ends at $t_{\mathrm{r}} = 1800$ s. This creates a 300 s pre-closure management interval and a 600 s active protection interval. These timings are scenario-design parameters and are not derived from aircraft performance, descent dynamics, or emergency-response procedures.
The experiment is deterministic. Multiple random seeds, stochastic demand perturbations, and behavioural parameter variations are not evaluated. The results should therefore be interpreted as a controlled comparative assessment of AGTM strategies rather than a statistical estimate of real-world performance.
Figure 4 shows the geographic study area. Figure 5 shows the corresponding SUMO representation.



Figure 4 – Map view of the selected Munich study area used for the simulation-based emergency landing-site assessment.



Figure 5 – SUMO representation of the selected Munich road network, including arterial corridors, local links, and signalised intersections used for AGTM-controlled emergency landing-site assessment.

3.2 Representation of the Emergency Landing Request

The AAM emergency is modelled as a time-critical air–ground traffic-management event. The selected signalised intersection is treated as a temporary candidate area that must be protected from vehicle entry during the landing-safety window. The surrounding road network represents the ground-traffic system affected by this restriction.

The model does not simulate aircraft descent, touchdown dynamics, downwash, obstacle clearance, surface suitability, or certification requirements. Instead, it isolates the ground-control problem: once a signalised intersection is selected as a candidate emergency landing area, how does the road network respond when that intersection is temporarily restricted?

The emergency event consists of four phases:

a. **Normal operation:** traffic follows the original SUMO signal programmes until the AGTM decision time t_{d}.
b. **Pre-closure management:** during $(t_{\mathrm{d}}, t_{\mathrm{c}})$, coordinated scenarios may reroute vehicles away from the candidate landing intersection.
c. **Active protection:** during $(t_{\mathrm{c}}, t_{\mathrm{r}})$, the selected traffic-light system is forced to an all-red state to restrict vehicle access.
d. **Recovery:** after t_{r}, the original signal programme is restored and post-event network recovery

is observed.

The $600\ \mathrm{s}$ active protection interval is a controlled stress-test window for evaluating traffic response to temporary intersection protection. It does not represent the complete real-world incident duration, which may include passenger evacuation, emergency-service access, investigation, aircraft removal, and road reopening procedures.

3.3 AGTM Control Logic

AGTM is implemented through three runtime control actions: signal restriction, route intervention, and candidate-site selection. The implementation is based on TraCI runtime control scripts that retrieve traffic-light states, controlled lanes, vehicle routes, vehicle speeds, halted vehicles, and waiting-time indicators at each simulation step. Rerouting interventions are applied only to vehicles whose planned routes intersect the selected candidate area.

Signal restriction is implemented by storing the original signal programme of the selected traffic-light system and overriding the active state with an all-red phase during $(t_\mathrm{c}, t_\mathrm{r})$. At t_r, the original programme is restored. The all-red override represents the emergency restriction of vehicle access to the candidate landing area. And, the all-red state is used as a simplified representation of emergency access restriction, not as a deployable signal-control plan.

The implementation does not include a staged green-clearance phase before closure. Vehicles already close to or inside the junction are therefore not explicitly flushed before the all-red protection begins. This makes the closure response conservative from a traffic-disruption perspective. A deployable AGTM system would require a staged sequence that first clears the conflict area and then locks all approaches.

Route intervention is applied in coordinated scenarios during the pre-closure interval. Vehicles whose planned routes include the selected intersection or its controlled approaches are identified through TraCI. Where feasible, SUMO rerouting is used to redirect these vehicles before the active protection interval begins. This intervention is designed to reduce vehicle accumulation near the landing candidate before access is restricted.

Candidate-site selection is applied in the dynamic scenarios. At t_d, all active signalised intersections are evaluated using live traffic states. Two selection logics are tested. S3 selects the lowest-disruption candidate based on current traffic conditions. S4 first filters candidates using simplified AAM-oriented controllability proxies and then selects the best candidate among the filtered set.

3.4 Scenario Design

Five scenarios are defined to isolate the effect of increasing AGTM coordination:

- **S0: Baseline operation.** The network operates normally for the full simulation horizon without emergency intervention.
- **S1: Uncoordinated all-red closure.** A fixed signalised intersection is forced to all-red during $(t_\mathrm{c}, t_\mathrm{r})$ without advance rerouting or site selection.
- **S2: Fixed-site coordinated response.** The same fixed intersection is used as in S1, but vehicles approaching or routed through the selected area are rerouted during $(t_\mathrm{d}, t_\mathrm{c})$ before the all-red protection begins.
- **S3: Traffic-aware dynamic selection.** All signalised intersections are evaluated at t_d. The selected site is the intersection with the lowest traffic-disruption cost according to Eq. (1). This scenario represents a traffic-management strategy that favours locations with lower immediate network impact.
- **S4: AAM-constrained dynamic selection.** Candidate intersections are first filtered using minimum controlled-lane and signal-state thresholds. The remaining intersections are ranked according to Eq. (4). This scenario represents a selection strategy that gives greater priority to larger and more controllable signalised intersections.

The scenario sequence is designed to separate four effects. S1 measures the disruption caused by an uncoordinated emergency closure. S2 tests whether pre-closure rerouting reduces this disruption. S3 tests whether real-time traffic-aware site selection improves performance relative to a fixed site. S4 tests the effect of adding AAM-oriented controllability constraints to the dynamic selection process.

3.5 Candidate-Selection Model

Let TLS denote the set of active traffic-light systems in the SUMO network. Each signalised intersection $i \in \mathrm{TLS}$ is treated as a potential candidate in the dynamic scenarios. Candidate evaluation is performed at the AGTM decision time t_d.

For each candidate i, four indicators are extracted:

- $N_i(t_\mathrm{d})$: number of vehicles on controlled approaches,
- $H_i(t_\mathrm{d})$: number of halted vehicles on controlled approaches,
- $R_i(t_\mathrm{d})$: number of vehicles whose planned route conflicts with the candidate intersection,
- L_i: number of controlled lanes.

In S3, the objective is to minimise immediate traffic disruption. The candidate cost is defined as:

$$C_i^{\mathrm{S3}} = 1.0\,N_i(t_\mathrm{d}) + 2.0\,H_i(t_\mathrm{d}) + 0.5\,R_i(t_\mathrm{d}) + 0.2\,L_i. \tag{1}$$

The selected S3 candidate is:

$$i_{\mathrm{S3}}^{*} = \arg\min_{i \in \mathrm{TLS}} C_i^{\mathrm{S3}}. \tag{2}$$

The positive lane term in Eq. (1) penalises larger intersections because their restriction is expected to affect more traffic movements. S3 therefore tends to favour smaller or less disruptive intersections.

In S4, the candidate set is first restricted using two simplified controllability proxies. Let P_i denote the signal-state length of intersection i. The feasible AAM-constrained candidate set is:

$$\mathrm{TLS}_{\mathrm{AAM}} = \{\, i \in \mathrm{TLS} \mid L_i \geq 8,\ P_i \geq 12 \,\}. \tag{3}$$

The thresholds $L_i \geq 8$ and $P_i \geq 12$ are heuristic filters. They do not certify physical landing feasibility. They are used only to exclude small or weakly controlled junctions from the AAM-constrained candidate set. Controlled-lane count is interpreted as a proxy for traffic-control authority and intersection scale. Signal-state length is interpreted as a proxy for signal-control complexity.

For candidates that satisfy Eq. (3), the S4 cost is:

$$C_i^{\mathrm{S4}} = 1.0\,N_i(t_\mathrm{d}) + 2.0\,H_i(t_\mathrm{d}) + 0.5\,R_i(t_\mathrm{d}) - 0.5\,L_i. \tag{4}$$

The selected S4 candidate is:

$$i_{\mathrm{S4}}^{*} = \arg\min_{i \in \mathrm{TLS}_{\mathrm{AAM}}} C_i^{\mathrm{S4}}. \tag{5}$$

The scoring functions are not intended to represent calibrated suitability models. They are rule-based ranking functions used to compare two operational selection behaviours under identical simulation conditions. S3 favours candidates with lower immediate traffic disruption, while S4 introduces AAM-oriented controllability constraints by excluding small or weakly controlled intersections and rewarding larger controlled-lane coverage.

The negative lane term in Eq. (4) rewards larger controlled-lane coverage. This reflects the assumption that larger signalised intersections may offer stronger traffic-control authority and may be more relevant for ground-side AAM emergency preparation than small junctions. This assumption is operational, not aeronautical.

The coefficients in Eqs. (1) and (4) are fixed heuristic parameters. They are not calibrated, normalised, or transferable suitability scores. Their purpose is to create two reproducible selection behaviours within the same network: S3 prioritises minimum traffic disruption, while S4 introduces a controllability-oriented AAM constraint. Future work will normalise the indicators and calibrate the weights against site geometry, aircraft requirements, traffic demand, emergency-response policy, and safety constraints. Table 1 summaries the scenario differences.

Table 1 – Summary of evaluated AGTM control scenarios.

Scenario	Site selection	Rerouting	Filtering	Lane term
S0	None	No	No	None
S1	Fixed	No	No	None
S2	Fixed	Yes	No	None
S3	Dynamic, Eq. (1)	Yes	No	Penalty, $+0.2L_i$
S4	Dynamic, Eq. (4)	Yes	Yes, Eq. (3)	Reward, $-0.5L_i$

3.6 Performance Indicators

Scenario performance is evaluated using network-level and local indicators.The selected indicators are intended to capture both the ability to protect the candidate landing area locally and the secondary disruption imposed on the wider road network.

Network-level indicators measure the wider disruption caused by temporary intersection protection:

- average network speed,
- number of halted vehicles,
- accumulated waiting time,
- active vehicle count,
- post-release recovery behaviour.

Local indicators measure the immediate condition around the selected landing candidate:

- vehicle count on controlled approaches,
- halted vehicles on controlled approaches,
- route conflicts involving the selected intersection,
- number of rerouted vehicles.

Accumulated waiting time is interpreted as a cumulative network-level delay indicator. When reported for a specific phase, it refers to the maximum cumulative waiting-time value observed within that phase, not to delay generated exclusively during that phase.

The methodology therefore evaluates the emergency landing request as a ground-side traffic-management problem. It compares uncoordinated closure, fixed-site coordination, traffic-aware dynamic selection, and AAM-constrained dynamic selection under identical baseline microscopic traffic conditions.

4. Results and Analysis

The results evaluate whether AGTM intervention can reduce the traffic disruption caused by temporary restriction of a signalised intersection for emergency AAM safety landing-site preparation. The analysis focuses on four effects: network-level disruption, local landing-site clearance, rerouting burden, and post-release recovery. The AGTM decision is triggered at $t_{\mathrm{d}} = 900\ \mathrm{s}$, the protection interval begins at $t_{\mathrm{c}} = 1200\ \mathrm{s}$, and the selected intersection is released at $t_{\mathrm{r}} = 1800\ \mathrm{s}$. All values are deterministic scenario estimates under identical network, demand, and timing assumptions.

4.1 Network-Level Response

The first result is that uncoordinated emergency restriction produces a clear network-level penalty. Compared with the baseline S0, the uncoordinated all-red closure in S1 reduces the protection-phase mean speed from $3.255\ \mathrm{m/s}$ to $2.618\ \mathrm{m/s}$ and increases the maximum number of halted, vehicles from 774 to $1{,}033$ (Table 2). This confirms that temporary intersection protection cannot be treated as a purely local control action; the restriction propagates into the surrounding road network.

The baseline itself is not free-flowing. During the recovery phase, S0 reaches $1{,}874$ halted vehicles and a mean speed of only $1.880\ \mathrm{m/s}$. Therefore, comparisons with S0 are used to characterise the background congestion level, not to claim that emergency intervention improves normal operation.

The waiting-time column reports the maximum network total waiting-time value observed within the corresponding phase. It should not be interpreted as incremental delay generated exclusively during that phase.

The second result is that fixed-site rerouting alone provides only partial mitigation. S2 reduces the protection-phase maximum halted vehicles from 1033 to 978, corresponding to a 5.32% reduction relative to S1. However, the maximum network total waiting-time value increases from $95091\ \mathrm{s}$ in S1 to $96706\ \mathrm{s}$ in S2. This indicates that rerouting can reduce part of the queue associated with the

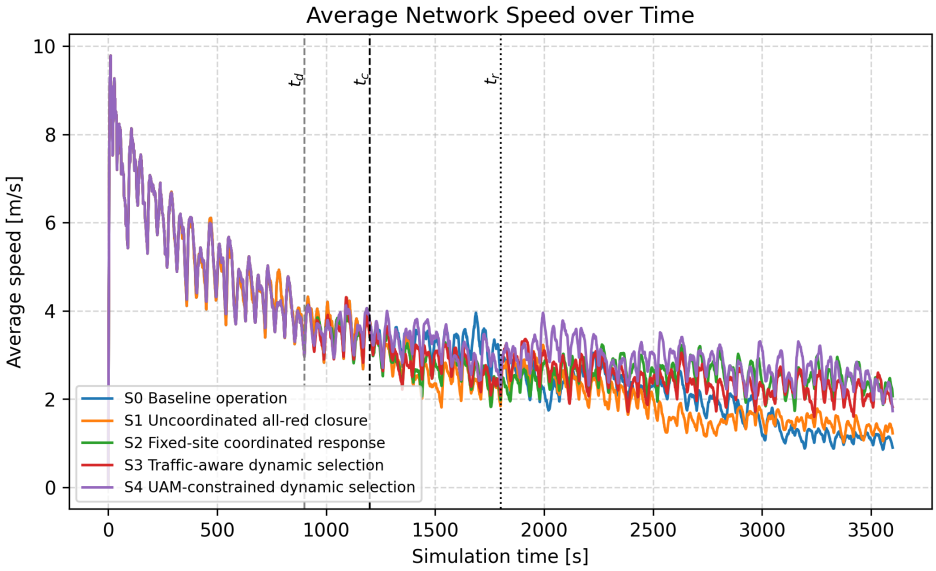


Figure 6 – Average network speed over time for the evaluated AGTM scenarios. The vertical markers indicate the AGTM decision time t_{d}, closure start t_{c}, and release time t_{r}.

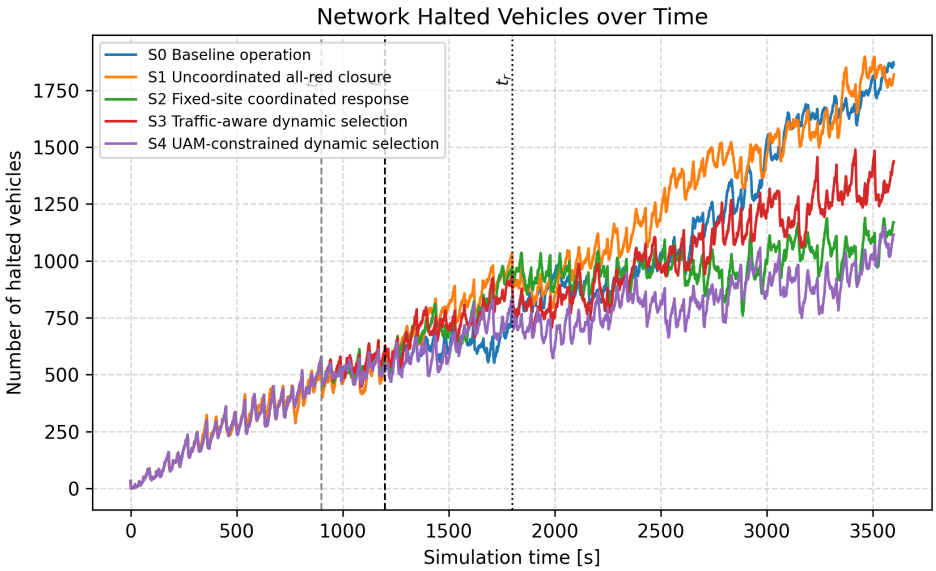


Figure 7 – Network-wide halted vehicles over time for the evaluated AGTM scenarios.

restricted intersection, but may transfer delay to alternative corridors when the landing site remains fixed.

Dynamic site selection produces a stronger network response. S3 reduces the protection-phase maximum halted vehicles to 937, while S4 further reduces this value to 837. Relative to S1, these correspond to reductions of 9.29% and 18.97%, respectively. S4 also achieves the highest protection-phase mean speed, $3.203\ \mathrm{m/s}$, indicating the strongest network-level mitigation during the active protection interval.

4.2 Recovery Behaviour

The recovery phase indicates whether AGTM intervention only displaces congestion or also reduces residual network burden after the restriction is released. S4 provides the strongest recovery performance, with a mean speed of $2.820\ \mathrm{m/s}$, $1,151$ halted vehicles, and a maximum network total waiting-time value of $77,019\ \mathrm{s}$. In contrast, S3 reaches $2.414\ \mathrm{m/s}$, $1,489$ halted vehicles, and $89,412\ \mathrm{s}$. Thus, although S3 is stronger for local landing-site clearance, S4 produces lower residual network disruption after release.

Table 2 – Phase-based network performance during protection and recovery intervals.

Scenario	Protection phase			Recovery phase		
	Speed [m/s]	Halted [veh]	Waiting [s]	Speed [m/s]	Halted [veh]	Waiting [s]
S0	3.255	774	52396	1.880	1874	236497
S1	2.618	1033	95091	1.815	1898	187684
S2	2.772	978	96706	2.530	1189	92026
S3	2.855	937	58449	2.414	1489	89412
S4	3.203	837	59462	2.820	1151	77019

Recovery time is evaluated using a baseline-relative first-crossing metric:

$$t_{\mathrm{rec}} = \min\{t \mid t \geq t_{\mathrm{r}},\ \bar{v}_{\mathrm{S}x}(t) \geq 0.95\,\bar{v}_{\mathrm{S0}}(t)\} - t_{\mathrm{r}}. \tag{6}$$

Using this definition, S4 reaches the recovery threshold fastest after 5 s, followed by S3 after 8 s, S1 after 12 s, and S2 after 68 s. These values should be interpreted as first-crossing indicators, not as proof of sustained stable recovery, because the threshold is defined relative to the time-varying and already congested baseline trajectory. The longer recovery time of S2 supports the interpretation that fixed-site rerouting can create secondary congestion that dissipates more slowly than in the dynamic-selection scenarios.

4.3 Local Landing-Site Clearance

The local queue results show a different ranking from the network-level indicators. S1 produces the largest local queue, reaching 111 halted vehicles near the selected landing intersection during the protection interval. S2 reduces this value only slightly to 103, despite rerouting 614 vehicles. This confirms that rerouting around a fixed high-demand landing site is not sufficient to guarantee local clearance.

S3 achieves the strongest local clearance. The maximum local halted vehicles decrease to 27, corresponding to a 75.68% reduction relative to S1. This is consistent with the S3 scoring logic, which favours low-disruption intersections from the full candidate set. S4 produces a higher local maximum of 58 halted vehicles, but still reduces the local queue by 47.75% relative to S1. The higher local queue in S4 is expected because this scenario excludes smaller candidates and favours larger signalised intersections according to the selected ground-side controllability proxies.

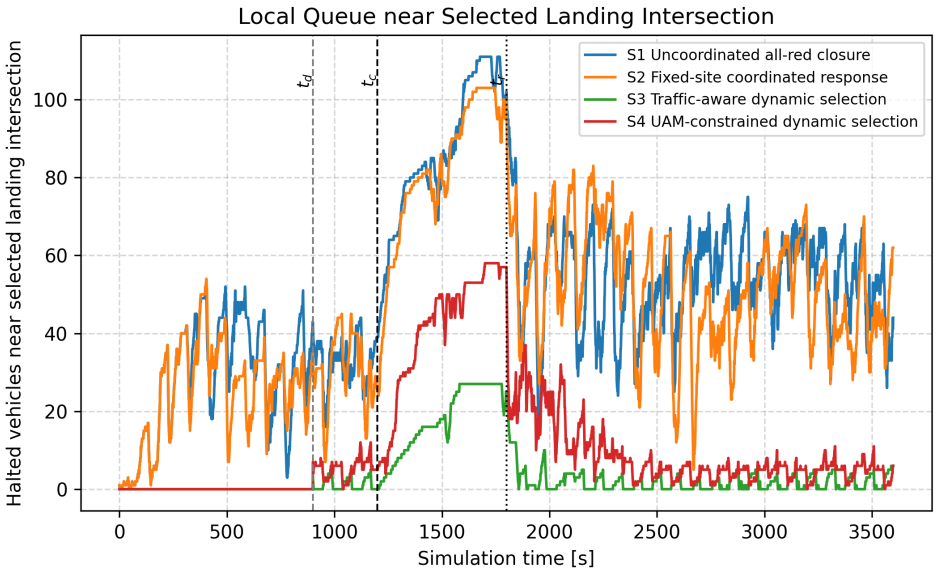


Figure 8 – Local halted vehicles near the selected landing intersection. S3 provides the strongest local clearance, while S4 accepts a higher local queue in exchange for stronger network-level resilience.

4.4 Trade-off Between Local Clearance and Network Resilience

The improvement index $\Omega_{\%}$ is used to compare coordinated AGTM strategies against the uncoordinated emergency closure S1 during the protection interval:

$$\Omega_{\%} = \frac{\mathscr{M}_{\mathrm{S1}} - \mathscr{M}_{\mathrm{S}x}}{\mathscr{M}_{\mathrm{S1}}} \times 100. \tag{7}$$

where $\mathscr{M}_{\mathrm{S1}}$ is the metric value under S1 and $\mathscr{M}_{\mathrm{S}x}$ is the corresponding value for the evaluated strategy. Positive values indicate improvement for metrics where lower values are preferred.

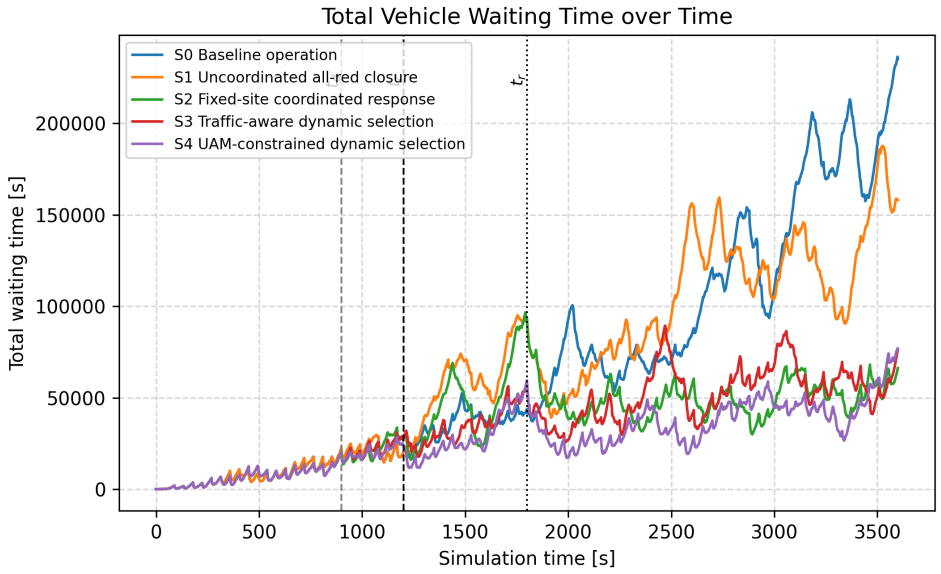


Figure 9 – Network total waiting time over time for the evaluated scenarios.

Table 3 – Relative improvement of coordinated AGTM strategies with respect to S1 during the protection interval.

Scenario	Max halted [veh]	Imp. [%]	Max waiting [s]	Imp. [%]	Max local halted [veh]	Imp. [%]
S1	1 033	0.00	95 091	0.00	111	0.00
S2	978	5.32	96 706	-1.70	103	7.21
S3	937	9.29	58 449	38.53	27	75.68
S4	837	18.97	59 462	37.47	58	47.75

Table 3 shows that S3 and S4 optimise different objectives. S3 provides the strongest local clearance, reducing local halted vehicles by 75.68% relative to S1. S4 provides the strongest network-level reduction in halted vehicles, reducing the protection-phase maximum by 18.97%. This difference reflects the selection logic rather than an inconsistency in the results. S3 favours the least disruptive candidate, while S4 favours candidates that are larger and more controllable according to the selected ground-side proxy variables.

The selected-site comparison in Table 4 explains this trade-off. S3 reroutes only 87 vehicles and reaches the lowest local queue of 27 halted vehicles. S4 reroutes 341 vehicles and accepts a higher local queue of 58 halted vehicles. However, S4 reduces the recovery-phase maximum network total waiting-time value from $89\,412$ s in S3 to $77\,019$ s. This corresponds to a reduction of $12\,393$ s, or approximately 13.86%, relative to S3 during recovery.

Table 4 – Selected-site comparison for dynamic AGTM scenarios.

Scenario	Selection score	Rerouted vehicles	Max local halted
S3	36.8	87	27
S4	141.5	341	58

The selected-site results show that the lowest-disruption traffic candidate is not necessarily the most operationally relevant candidate for AAM emergency landing-site preparation. S3 is better for imme-

diate local clearance. S4 is better for network-level resilience because it selects a larger controlled intersection and redistributes more traffic before the protection interval.

4.5 Key Findings

First, uncoordinated all-red closure creates both local and network-level disruption. This confirms that emergency landing-site protection cannot be treated as an isolated intersection-control problem.
Second, fixed-site rerouting provides limited benefit. S2 reduces halted vehicles relative to S1, but increases protection-phase waiting time and has the longest recovery threshold time. This indicates that rerouting without site selection may displace congestion rather than remove it.
Third, traffic-aware dynamic site selection is most effective for local clearance. S3 achieves the lowest local queue because it selects a low-disruption candidate from the full signalised-intersection set.
Fourth, AAM-constrained dynamic selection provides the strongest network-level compromise. S4 accepts a higher local queue than S3, but it produces the best protection-phase network speed, the lowest protection-phase halted-vehicle count, and the lowest recovery-phase waiting-time value among the emergency scenarios.
The main implication is that AGTM-supported AAM emergency landing management requires a trade-off between two objectives: clearing the selected landing-site approaches and preserving network-level resilience. A strategy that minimises local traffic disruption may not provide the best recovery behaviour. Conversely, a strategy that favours larger and more controllable signalised intersections may create higher local queues but reduce residual network congestion after the emergency restriction is released.

5. Discussion and Limitations

5.1 Discussion

The results show that emergency landing-site preparation at signalised intersections is not a local signal-control problem alone. It is an air–ground traffic-management problem in which the selected landing candidate affects local clearance, rerouting load, queue propagation, and network recovery. This is the central operational finding of the study.
The uncoordinated closure scenario S1 establishes the baseline failure mode. Setting the selected intersection to an all-red state can restrict vehicle entry, but it does not manage the traffic consequences of that restriction. The resulting increase in local and network-level congestion shows that intersection protection alone is insufficient for AAM emergency landing-site preparation. A protected ground area must therefore be produced through coordinated Advanced Ground Traffic Management (AGTM), not through isolated signal override.
The fixed-site coordinated response S2 demonstrates a second failure mode. Rerouting vehicles before the closure reduces the maximum number of halted vehicles relative to S1, but it does not provide the best overall mitigation. During the protection interval, S2 produces a maximum network total waiting-time value of $96{,}706\ \mathrm{s}$ and reaches the baseline-relative recovery threshold only after $68\ \mathrm{s}$. This indicates that rerouting around a fixed landing candidate can shift congestion into alternative corridors rather than remove it. The result supports the need to couple rerouting with landing-site selection instead of treating the landing site as a predetermined input.
The dynamic strategies clarify the contribution of site selection. S3 provides the strongest local landing-site clearance, reducing the maximum number of local halted vehicles to 27, compared with 111 in S1. This shows that traffic-aware dynamic selection can identify candidates that are easier to clear under the current traffic state. However, S3 achieves this by favouring smaller, lower-disruption signalised junctions. This improves the local clearance metric, but limits the scale of network-wide demand redistribution and is associated with weaker recovery-phase performance than S4.
S4 changes the selection objective. By filtering candidate intersections and rewarding larger controlled-lane coverage, S4 shifts the problem from pure traffic-disruption minimisation toward controllability-oriented candidate selection. This produces a higher local queue than S3, with 58 local halted vehicles instead of 27, but improves the network-level response. S4 achieves the lowest protection-phase maximum halted-vehicle count among the emergency scenarios, 837, and the strongest recovery-phase performance, with a mean speed of $2.820\ \mathrm{m/s}$ and $1{,}151$ halted vehicles. This shows that the lowest-disruption traffic candidate is not necessarily the most operationally useful candidate for AAM emergency landing-site preparation.

The main interpretation is therefore not that one strategy is universally best. S3 and S4 optimise different operational priorities. S3 is preferable when the objective is to minimise immediate local disruption around the selected site. S4 is preferable when the objective is to balance local clearance with broader network resilience. This distinction matters because AAM emergency landing management cannot be reduced to selecting the easiest intersection to clear. The selected site must also support traffic restriction, rerouting, and recovery at the network level.

The contribution of this study is to demonstrate a ground-side evaluation layer for AAM emergency landing-site preparation. Much of the reviewed AAM contingency literature focuses on airborne diversion, reachable landing alternatives, or vehicle-side emergency management. This study instead evaluates what happens at a signalised urban intersection when selected as a temporary ground-side safety candidate. It links candidate selection, signal restriction, rerouting, local queue clearance, and post-release recovery in one simulation framework. The results show that AGTM design choices directly change the operational feasibility of preparing a temporary landing area.

The findings should not be interpreted as evidence that signalised intersections are physically suitable landing sites. The study does not certify landing feasibility. It evaluates whether selected signalised intersections can be managed as controllable ground-side candidates under different AGTM strategies. Controlled-lane coverage and signal-state length are used only as simplified ground-side controllability proxies. They do not replace geometric, structural, pedestrian, obstacle-clearance, downwash, or aviation-safety assessment.

5.2 Limitations

First, the aerial vehicle dynamics are not modelled. The simulation does not include eVTOL performance, descent trajectory, battery state, degradation mode, glide-path feasibility, touchdown dynamics, or flight-control constraints. The emergency event is represented as a ground-side management request. Therefore, the results describe the traffic-management feasibility of preparing a selected candidate area, not the physical feasibility of landing an aircraft there.

Second, the candidate-selection models are heuristic. The S3 and S4 scoring functions are rule-based ranking functions used to create reproducible decision behaviour in the simulation. Their weights and thresholds are not calibrated, optimised, or transferable constants. They should be interpreted as scenario-specific implementation parameters for comparing two AGTM logics: minimum-disruption selection and controllability-oriented selection.

Third, the AAM-constrained filtering logic uses indirect proxies rather than certified landing-site criteria. Controlled-lane coverage and signal-state length approximate ground-side controllability and intersection scale. They do not evaluate touchdown area dimensions, surface quality, obstacle clearance, overhead infrastructure, street furniture, tram wires, pedestrian exposure, or approach and departure corridors. These factors must be included before any real-world landing-site suitability claim can be made.

Fourth, the road-traffic model excludes several operational complexities. Pedestrians, cyclists, public transport priority, emergency-service vehicles, communication delays, partial driver compliance, enforcement uncertainty, and behavioural adaptation are not included. The model assumes that rerouting commands are executed where route alternatives are available. This simplifies the real-world implementation problem.

Fifth, the evaluation is deterministic and limited to one Munich network configuration. The study does not include multiple random seeds, stochastic demand variation, sensitivity analysis of event timing, or calibration of scoring weights. The results should therefore be interpreted as a deterministic feasibility comparison of AGTM strategies in one case-study network, not as statistically generalisable performance estimates.

Sixth, the current study applies the emergency all-red restriction immediately at the start of the protection interval. A more realistic deployment would include a staged clearance sequence, in which vehicles already inside or immediately upstream of the selected junction are first allowed to clear the conflict area before the approaches are locked. Implementing such a clearance phase in TraCI may further reduce local halted vehicles and should be evaluated in future work.

Despite these limitations, the study provides a useful proof of concept. It shows that all-red closure alone is insufficient, fixed-site rerouting can displace congestion, and dynamic site selection changes both local and network-level outcomes. Most importantly, the comparison between S3 and S4 shows

that AAM emergency landing-site preparation requires a balance between local clearance, ground-side controllability, and network resilience.

6. Conclusion

This paper evaluated the ground-side traffic-management problem associated with using signalised urban intersections as temporary emergency safety landing-site candidates for Advanced Air Mobility vehicles. A microscopic SUMO/TraCI simulation framework was developed for an urban road network in Munich and used to compare five operational conditions: baseline operation, uncoordinated all-red closure, fixed-site coordinated rerouting, traffic-aware dynamic selection, and AAM-constrained dynamic selection.

The results show that temporary intersection protection cannot be treated as a simple traffic-light override. The uncoordinated all-red closure creates substantial disruption, while fixed-site rerouting provides only partial mitigation and can transfer delay to alternative corridors. This demonstrates that emergency landing-site preparation requires coordinated AGTM that considers both the selected intersection and the surrounding road network.

The main finding is a trade-off between local clearance and network resilience. S3 achieves the strongest local landing-site clearance by selecting lower-disruption candidates under current traffic conditions. S4 accepts a higher local queue, but selects larger signalised intersections according to ground-side controllability proxies and achieves stronger network-level recovery performance. Therefore, the lowest-disruption traffic candidate is not necessarily the most operationally meaningful candidate for AAM emergency landing-site preparation.

The exact contribution of this study is a simulation-based ground-side evaluation framework for AGTM-supported AAM emergency landing management. The framework connects dynamic candidate selection, signal restriction, rerouting, local clearance, and post-release recovery. It does not certify signalised intersections as physically suitable landing sites. Instead, it shows how selected intersections can be evaluated as controllable ground-side candidates within an integrated air–ground traffic-management process.

Future work should extend the framework in four directions. First, physical landing-site constraints should be added, including landing geometry, obstacle clearance, surface quality, downwash effects, and approach/departure corridors. Second, multimodal road users such as pedestrians, cyclists, public transport, and emergency vehicles should be included. Third, operational uncertainty should be modelled through communication delays, partial driver compliance, stochastic demand variation, and multiple simulation seeds. Fourth, the S3 and S4 scoring parameters should be calibrated and tested through sensitivity analysis before the approach is applied to other urban networks.

7. Acknowledgement

This work was funded by the Bavarian Ministry of Economic Affairs, Regional Development and Energy through the URTEAM project (HAM-2208-0024).

8. Contact Author Email Address

Pavithra Sathya Kumar pavithra.sathya@unibw.de

9. Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] ICAO. Extended diversion time operations (EDTO) manual (Doc 10085), 2017. Montréal, Quebec, Canada.

[2] ICAO. *Annex 6 - Operation of Aircraft - Part I - International Commercial Air Transport - Aeroplanes*. ICAO, 12th edition, July 2022.

[3] FAA/ NASA. Urban air mobility (UAM) concept of operations v2.0. Technical report, NASA/ FAA, 2023.

[4] EUROCONTROL Consortium. U-space ConOps and Architecture (Edition 4). Technical report, CORUS-XUAM Project, 2023.

[5] M Schultz et al. Intermodal safety net for urban air mobility. In *TRC-30 Conference*, Heraklion, Greece, 2024.

[6] Pavithra Sathya Kumar, Ehsan Asadi, and Michael Schultz. Multimodal traffic coordination for safety landings. In *Proceedings of the First US-Europe Air Transportation Research and Development Symposium (ATRDS 2025)*, 2025.

[7] Irene M Gregory, Newton H Campbell, Natasha A Neogi, Jon B Holbrook, Jared A Grauer, Barton J Bacon, Patrick C Murphy, Daniel D Moerder, Benjamin M Simmons, Michael J Acheson, et al. Intelligent contingency management for urban air mobility. In *International Conference on Dynamic Data Driven Application Systems*, pages 22–26. Springer, 2020.

[8] Arinc Tutku Altun, Yan Xu, Gokhan Inalhan, Ignacio Vidal-Franco, and Michael Hardt. Contingency management concept generation for u-space system. In *2022 Integrated Communication, Navigation and Surveillance Conference (ICNS)*, pages 1–12. IEEE, 2022.

[9] Zahra Baleghi and Others. Realtime landing management for evtols. *Journal of Advanced Air Mobility*, 12:123–134, 2023.

[10] Euclides Carlos Pinto Neto, Derick Baum, Jorge Rady de Almeida Jr, João Batista Camargo Jr, and Paulo Sergio Cugnasca. Towards planning urban air mobility (uam) landing trajectories in emergencies. *Journal of Intelligent & Robotic Systems*, 111(1):22, 2025.

[11] Liqun Ma, Yangchen Ding, and Jianzhong Yang. Autonomous emergency gliding landing guidance and control of tilt-wing electric vertical take-off and landing for urban air mobility missions using control barrier functions. *Aerospace*, 12(1):63, 2025.

[12] Bingze Xia, Iraj Mantegh, and Wenfang Xie. Integrated emergency self-landing method for autonomous uas in urban aerial mobility. In *2021 21st international conference on control, automation and systems (iccas)*, pages 275–282. IEEE, 2021.

[13] Md Shah Alam and Jared Oluoch. A survey of safe landing zone detection techniques for autonomous unmanned aerial vehicles (uavs). *Expert Systems with Applications*, 179:115091, 2021.

[14] Felix Wachter, Samuel Würtz, and Andreas Rupp. An analytical investigation on the use of passenger drones in the rendezvouz-system of emergency medical services for landing site identification. In *2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1449–1455. IEEE, 2024.

[15] Karima Benhamza and Hamid Seridi. Adaptive traffic signal control in multiple intersections network. *Journal of Intelligent & Fuzzy Systems*, 28(6):2557–2567, 2015.

[16] Mohamed Essa and Tarek Sayed. Self-learning adaptive traffic signal control for real-time safety optimization. *Accident Analysis & Prevention*, 146:105713, 2020.

[17] Malose Mathiane, Chunling Tu, Pius Adewale Owola, and Mukatshung Claude Nawej. A sumo simulation study on vanet-based adaptive traffic light control system. In *International conference on advances in electrical and computer technologies*, pages 225–237. Springer, 2021.

[18] Mark Colley, Julian Czymmeck, Pascal Jansen, Luca-Maxim Meinhardt, Patrick Ebel, and Enrico Rukzio. Uam-sumo: Simulacra of urban air mobility using sumo to study large-scale effects. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 1000–1004. IEEE, 2025.

[19] Selim Reza, Marta Campos Ferreira, José JM Machado, and Joao Manuel RS Tavares. A citywide td-learning based intelligent traffic signal control for autonomous vehicles: Performance evaluation using sumo. *Expert Systems*, 42(1):e13301, 2025.

[20] Abdelgwad Elashry, Ahmed Elasmar, Yasmeen Khalifa, Mohamed Bahnassy, and Alyaa A Hamza. Mitigating ddos threats in vanets: A real-time simulation approach using sumo and ns-3. In *2025 7th Novel Intelligent and Leading Emerging Sciences Conference (NILES)*, pages 254–258. IEEE, 2025.

[21] Jiangchen Li, Chen Qiu, Liqun Peng, and Tony Z Qiu. Signal priority request delay modeling and mitigation for emergency vehicles in connected vehicle environment. *Transportation research record*, 2672(18):45–57, 2018.

[22] Chun Ting Wu, Yung-Jen Hsu, Chia-Heng Tu, and Shaohua Wang. Traffic signal preemption and prioritization for high-density road networks. In *Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing*, pages 410–419, 2024.