multi-Risk sciEnce for resilienT commUnities undeR a changiNg climate

Codice progetto MUR: **PE00000005** – CUP LEAD PARTNER: I33C22006910006

Deliverable title: **Report on intra- and inter-CI service impact**

Deliverable ID: **DV6.2.4**

Due date: **30/08/2025**

Submission date: **28/11/2025**

AUTHORS

Luisa Lavalle (ENEA), Daniele Fava (ENEA), Zahra Lahijanian (ENEA), Pierluigi Claps (PoliTO)

Approved by Luisa Lavalle (WP leader)

Technical references

Project Acronym	RETURN
Project Title	multi-Risk sciEnce for resilienT commUnities undeR a changiNg climate
Project Coordinator	Domenico Calcaterra UNIVERSITA DEGLI STUDI DI NAPOLI FEDERICO II domcalca@unina.it
Project Duration	December 2022 – November 2025 (36 months)

Deliverable No.	DV6.2.4
Dissemination level*	PU
Work Package	WP2 - Asset systems definition and characterization
Task	T2.4 – Analysis of service continuity and classification of the intra- and inter-CI service failures
Lead beneficiary	ENEA
Contributing beneficiary/ies	

* PU = Public

PP = Restricted to other programme participants (including the Commission Services)

RE = Restricted to a group specified by the consortium (including the Commission Services)

CO = Confidential, only for members of the consortium (including the Commission Services)

Document history

Version	Date	Lead contributor	Description
0.1	xx.xx.2025	Luisa Lavalle (ENEA)	First Draft
0.2	xx.xx.2025	Luisa Lavalle (ENEA)	Critical review and proofreading
0.3	xx.xx.2025		Edits for approval
1.0	xx.xx.2025		Final version

ABSTRACT

Building on a simulation framework developed and described in previous RETURN deliverables, this work demonstrates one of its possible applications at scale to real-world infrastructure systems, highlighting its capability to uncover potential critical components, non-intuitive vulnerabilities, and cascading effects across drinking water distribution systems, railway networks, and road traffic systems. Developed within the RETURN project, this work advances the state of the art in infrastructure resilience assessment by combining exhaustive failure-oriented simulations with domain-specific, physics-based, and operationally realistic models.

The core methodological innovation lies in the systematic use of extensive or near-exhaustive simulation campaigns—covering tens of thousands of failure scenarios—to move beyond expert-driven, asset-by-asset risk assessments. By adopting an all-hazards, hazard-agnostic screening approach, the analyses deliberately decouple the identification of critical components from specific triggering events, allowing "vulnerabilities" to be detected based solely on their functional and systemic impact. Hazard-specific analyses are then reserved for a reduced subset of components where they are most likely to influence system performance, enabling a more efficient and proportionate allocation of modelling and mitigation efforts.

In the water domain, more than 18,000 pipe-break simulations and multiple pump-failure scenarios reveal that approximately 10–15% of components account for the majority of service disruptions, while most failures produce limited effects. The analyses highlight pressure loss, breakflow, and energy dependence as dominant drivers of vulnerability, with power blackouts emerging as a critical systemic risk. In the railway domain, over 1,100 simulations demonstrate how single-edge failures can propagate through timetable constraints and network topology, producing highly non-intuitive service disruptions that would be extremely difficult to anticipate without exhaustive simulation. A novel graph-based workaround enables full automation of failure injection in commercial railway simulation software, overcoming limitations in existing APIs and significantly extending their applicability to resilience studies. In the road domain, mesoscopic traffic simulations informed by Floating Car Data illustrate how bridge failures generate network-wide congestion patterns and rerouting dynamics, directly affecting accessibility to essential services.

Across all domains, results consistently show that a small number of components and locations dominate system-level risk, validating the application of the As Low As Reasonably Practicable (ALARP) principle as a practical decision framework. The work demonstrates that while human expertise remains indispensable for interpreting critical cases and defining mitigation strategies, systematic, machine-based simulation can considerably speed the process to reliably identify which scenarios truly matter. The integration of simulation outputs with GIS-based representations further enables direct adoption within existing decision-support systems, without requiring operators to interact with specialised simulation tools.

Overall, this deliverable provides a transferable, scalable, and operationally grounded framework for resilience assessment of critical infrastructures. It advances current practice by enabling evidence-based prioritisation, uncovering non-intuitive vulnerabilities, and supporting proportionate, risk-informed decision-making for disaster risk management and infrastructure protection.

Table of contents

Technical references 3
Document history 4
ABSTRACT 5
Table of contents 6
List of Tables 7
List of Figures 7
1. Introduction 9
1.1 Scope and Relation to other project activities 9
1.2 Document Structure 9
2. Drinking Water Distribution Network 10
2.1 Methodology 10
2.2 Simulations and Results 11
2.2.1 Pipe Rupture 12
2.2.2 Pump failure 17
2.3 Findings and Discussion 21
3. Railways 23
3.1 Methodology 23
3.2 Simulations and Results 26
3.2.1 Simulations overview 26
3.2.2 Results 28
3.3 Findings and Discussion 34
4. Roads 36
4.1 Methodology 37
4.1.1 Overview of simulation approach 37
4.1.2 Data collection and pre-processing 38
4.1.3 Model Development Steps 46
4.2 Simulations and Results 47
4.2.1 Prototype scenario design 47
4.2.2 Network connectivity tests 48
4.2.3 Model functionality tests (Designed scenarios) 52
4.3 Findings and Discussion 59
5. Conclusions 60
6. References 62

List of Tables

Table 1 – Summary of key indicators under total blackout 19
Table 2 – Modelling hazards in OpenTrack 24
Table 3 – Graph's attributes 42
Table 4 – Fundamental diagram values 43
Table 5 – O-D matrix of single platoon 48

List of Figures

Figure 1: Number of nodes with demand per pipe-break simulation 12
Figure 2: Number of isolated demand nodes per pipe-break simulation 12
Figure 3: Detailed statistics for the isolated-node cluster 13
Figure 4: Number of nodes with p≤0 per pipe-break simulation 13
Figure 5: Pressure distribution throughout nodes per pipe-break simulation 14
Figure 6: Distribution of disrupted demand severity across simulations 14
Figure 7: Percentage of demand affected by pressure drop 14
Figure 8: Number of isolated hydrants in pipe-break simulations 15
Figure 9: Number of hydrants with inadequate pressure 15
Figure 10: Percentage of unavailable hydrants 15
Figure 11: Breakflow per simulation 16
Figure 12: Overview of main KPIs 16
Figure 13: Demand nodes and Total demand across pump-failure scenarios 17
Figure 14: Not served demand nodes under single-pump failure 17
Figure 15: Pressure distribution across the network under single-pump failure 18
Figure 16: Demand with inadequate pressure under single-pump failure 18
Figure 17: Hydrants with inadequate pressure under single-pump failure 19
Figure 18: Demand nodes affected by service loss under single-pump emergency operation 20
Figure 19: Pressure distribution in demand nodes under single-pump emergency operation 20
Figure 20: Unserved demand for single-pump restart 20
Figure 21: Unavailable hydrants for single-pump restart 21
Figure 22: Simulation time 27
Figure 23: Simulation time and affected courses 27
Figure 24: Number of courses impacted and courses passing by blocked track 28
Figure 25: Number of courses impacted and courses passing by blocked track 28
Figure 26: Total arrived courses 30
Figure 27: Overview on high-speed courses 30
Figure 28: Overview on long-distance courses 31
Figure 29: Overview on freight courses 31
Figure 30: Overview on regional courses 32
Figure 31: Overview on deadhead courses 32
Figure 32: Courses arrived after re-routing evaluation 33
Figure 33: Workflow of the traffic modelling activity 36
Figure 34: Study area 37
Figure 35: Visualization of OD function of TomTom platform, a) Percentage distribution of trips by destination municipality with Pizzighettone as the origin, and b) Percentage distribution of trips by destination municipality with Pizzighettone as the destination 39
Figure 36: Visualization of traffic state function of TomTom platform 39
Figure 37: Link-Node graph of study area 40
Figure 38: Network, zoning and centroids visualizations 40
Figure 39: Zoning of study area 41
Figure 40: Centroids of zones 41
Figure 41: Extended network's speed limits 41

Figure 42: Fundamental diagram map 42
Figure 43: Flow-density relationship for each set of fundamental diagrams 43
Figure 44: Speed density relationship for each fundamental diagram 43
Figure 45: Assessment of traffic conditions along a specific route during the first week of August 2024 using TomTom data. 43
Figure 46: Desire lines representing the movements between two places, darker color shows higher number of displacements 44
Figure 47: Total number of trips counted for the wide network in 4 months (May, June, July, and October) 44
Figure 48: Total number of trips counted on each day of week for 4 months (May, June, July, and October) 45
Figure 49: Total number of trips in the study area by day of the week over four selected months (May, June, July, and October), showing the distribution of travel demand across weekdays and weekends. 45
Figure 50: Variation of total daily trips by day of the week from May to October, illustrating temporal changes in travel activity and weekday trends over the study period. 46
Figure 51: Traffic simulation framework 46
Figure 52: One pair of O-D for model functionality test 49
Figure 53: Test a. results 49
Figure 54: Pizzighettone bridges 50
Figure 55: Test b. results 50
Figure 56: Test c. results 51
Figure 57: External region trip counting concept 52
Figure 58: SP234 52
Figure 59: Origin-Destination matrix for hour 7 to 8 53
Figure 60: Speed demonstration based on real data from 7 to 8 in the morning 54
Figure 61: Simulated speed results for a typical Thursday of the day 54
Figure 62. Base scenario simulation O-D Seed 55
Figure 63: Comparison between TomTom traffic state data and simulated speeds on road segments approaching the Pizzighettone bridge. 55
Figure 64: Simulated speed 56
Figure 65: Area affected by closure 56
Figure 66: Simulated flow, scenario A-bridge closure 57
Figure 67. Simulate speed scenario A- Bridge closure 58

1. Introduction

1.1 Scope and Relation to other project activities

Within the RETURN project, the overarching goal of the Transversal Spoke TS2 (Spoke 6) is to strengthen Italian governance in disaster risk management by improving the understanding of environmental, natural, and anthropogenic risks affecting Critical Infrastructures.

Building on the requirements DW-RR1 and RL-RR1 developed in Task T6.2.1 and documented in Deliverable DV6.2.1, as well as on the modelling activities developed in Task T6.2.3, this deliverable follows up on Deliverable DV6.2.3 and – consistent with the commitments outlined in the Executive Working Plan – reports the results of the simulations performed on the water and railway networks. The objective was to identify and rank the most potential critical components based on the effects of single-component failures, using the Key Performance Indicators (KPIs) defined in Deliverable DV6.2.1.

This task required more time than initially anticipated due to the technical challenges encountered while executing thousands of simulations. Additional effort was also needed to collect the information and feedback required for the correct interpretation of the results. However, this process enabled the establishment of ongoing collaboration with infrastructure operators. Furthermore, as the project progressed, it was confirmed – fully in line with initial expectations – that these findings were instrumental in the construction of the scenarios developed within the ReturnVille and Proof of Concept Calabria activities. These scenarios, which are currently being elaborated into a few forthcoming publications, have enabled close collaboration with the hazards teams to assess which natural events could lead to the failure of the identified components, as well as to perform a preliminary cost–benefit analysis of the available mitigation measures.

Over the course of the project, the need to also model road traffic emerged, in order to assess the potential benefits of the proactive measures identified within the project and to identify the most critical bridges based on their impact on mobility. More generally, when dealing with natural hazards, a key strategic priority is ensuring continued accessibility to essential services, such as hospitals and emergency facilities. During events such as flooding—where parts of the transport network may become partially or fully inoperative—it is crucial to assess how mobility can be maintained and how evacuation operations can be effectively managed. This aspect is particularly relevant when evaluating the resilience of critical infrastructures under natural hazards and when analysing cascading effects triggered by consecutive events. For these reasons, a preliminary activity in the road traffic domain was also undertaken. However, due to limited time and personnel resources, this activity did not reach the same level of maturity as the analyses conducted for the other domains.

1.2 Document Structure

The report is organized as follows. Section 2 presents the analysis of the water distribution network simulations, while Section 3 summarizes the results for the railway domain. In both cases, the key functional aspects that require particular attention are highlighted. Section 4 provides a brief overview of the modelling and analysis activities conducted in the Pizzighettone bridge area concerning road traffic, which were undertaken in response to needs identified after the release of Deliverable DV6.2.3. Section 5 concludes the report with final remarks.

Before proceeding, it is important to note that, as this deliverable is public – while the case studies, unlike the virtual test beds, refer to real systems – the results presented here have been aggregated and anonymized, and the discussion is intentionally kept at a general level. Detailed results and per-simulation data have been archived and are available from the authors upon request.

2. Drinking Water Distribution Network

This section presents the analysis of the simulation results for the drinking water distribution network.

Section 2.1 provides a brief recap of the main methodological aspects – already detailed in Deliverable DV6.2.3 – for the reader's convenience. Section 2.2 describes the simulations performed and discusses the results, while Section 2.3 outlines the key findings.

2.1 Methodology

As discussed in Deliverable DV6.2.1 (Section 4.1.3) regarding the requirements analysis phase, and later in DV6.2.3, the study of functional vulnerabilities in water distribution networks greatly benefits from physics-based, large-scale, extensive simulation. Such an approach enables the quantification of impacts in both physical terms (flow losses, pressure anomalies) and functional terms (unserved users).

To ensure physically reliable outcomes, the simulation environment described in DV6.2.3 (Sections 2 and 3) integrated Infoworks WS PRO, a domain-specific water network simulator, and adapted the network model to reproduce the effects of natural events on infrastructure components. In line with the work program, this deliverable focuses on N-1 analyses. Beyond the extensive N-1 simulations reported here, the same environment can ingest hazard-related inputs (natural or anthropogenic), simulate any combination of failures, and propagate interdependency-driven effects to construct and analyse complex scenarios – as done in the ReturnLand/ReturnVille virtual test bed. These results – being developed into some papers currently in preparation – address scenarios such as earthquakes, debris flows, flooding, and electrical interdependency. Internal cascading effects are inherently captured by the domain simulator, whereas interdependency-induced effects are handled during the test-planning phase.

As described in DV6.2.1, interdependencies affecting water distribution networks may result in pipe breaks, loss of control, or pump shutdowns. In this deliverable, we adopt an all-hazards approach: the analysis is performed independently of the triggering hazard, and hazard scenarios are considered only at a later stage to determine which events may cause the identified failures. Given the context and time constraints, the simulations focused on pipe breaks (relevant to earthquakes and ground instabilities) and pump shutdowns (related to electrical dependencies and flooding). Under a cyber-attack scenario, by contrast, simulations would focus on valve or pump state changes or on forced pipe closures, all of which can also be represented within the simulation environment.

The analysed network includes:

- 18.688 pipes, for a total length of 834,4 Km;
- 28 wells;
- 54 pumping stations (28 well pumps and 26 flow-pressurization stations);
- 22.077 nodes, of which 6.405 have demand, with a total demand of approximately 345 l/s.

This corresponds to a medium-sized network serving approximately 200,000 inhabitants, representative of many moderately flat Italian cities without high-elevation reservoirs. While different network configurations may yield different insights, the same methodology is fully transferable. Moreover, Infoworks WS Pro is widely adopted by water utilities in Italy, including CAP, Hera, MM, Acea, Gori, SMAT, in addition to IREN. Collectively, these operators serve approximately 23,5 million residents across 13 regions.

The normal operating condition – validated during Task 2.3 – is defined as follows: all nodes are served, all pipes operate with positive flow and velocity, and the head at each demand node is at least 30 meters above ground elevation.

Before discussing the simulation outcomes, it is important to recall two modelling limitations – discussed in Deliverable DV6.2.3 – that influence the interpretation of the results:

1. No storage tanks.

   Storage tanks were not included because it was not possible to dimension them realistically, and doing so could have artificially increased or decreased the network's robustness. Their exclusion, however, also removes storage autonomy, which in real systems can mitigate or even fully compensate the effects of failures if the emergency is managed within the available autonomy window. At the same time, fragility curves (FEMA, 2024) indicate that tanks can easily go out of service – or even collapse – during natural events, meaning they may be unavailable or may even introduce additional hazards.

2. Constant demand.

   Demand was kept constant because information about real variability was not available, but this limitation affects only the quantification of the unserved demand metric. The number of affected demand nodes/users remains meaningful, and this simplification substantially reduced the simulation time: each test required simulating only one hour, since the 24-hour results would have been identical. Nevertheless, adjusting the results ex-post with an estimated weighting factor would not be difficult.

While these simplifications reduce the realism of some quantitative measures, they do not affect the identification of critical components, which remains robust across the tested configurations.

Finally, to ensure a consistent and meaningful assessment of network performance under failure conditions, the following operator-approved Key Performance Indicators (KPIs) were used:

- Pressure at nodes, with thresholds at $p \leq 0$, $p \leq 5$, $p \leq 15$, $p \leq 30$, as an indicator of service availability and quality. Pressure $p \leq 0$ indicates complete service outage, whereas the other thresholds represent reduced service quality – still potentially useful in emergency conditions.
- Unserved demand, used as an indicator of user inconvenience and as a proxy for characterizing the types of affected users.
- Isolated nodes.
- Breakflow (volume of water leaking from pipe breaks), which is relevant both for the hydraulic response and for potential cascading impacts on the environment and on geographically co-located infrastructures.
- Hydrant availability (pressure-based), which may be critical for emergency management involving fire response.

Although pipe velocity was computed by the simulator, it was not stored, as it is not part of the project KPIs.

2.2 Simulations and Results

A total of 18.688 N-1 pipe-break simulations and 109 pump-failure simulations were executed.

To preserve the confidentiality required when working with critical infrastructures, the results are anonymized and presented only in aggregated form.

Given that – according to the state-of-the-art fragility functions mentioned in Deliverable DV6.2.3 – earthquakes are expected to cause only a limited number of pipe breaks (and the most critical components had already been identified), the N-1 approach is generally adequate. For pumps, however, additional scenarios were analysed: a total power blackout (plausible in the event of failures in the electrical transmission network) and single-pump emergency restart.

Each simulation (1-hour duration) required 70-110 seconds and generated approximately 178 GB of output data. All simulations are fully reproducible, and their time evolution can be visualized by the operators. To avoid non-convergence issues, the maximum number of iterations was set to 999, while the default tolerance (1 l/s) was retained.

The following subsections summarize the main outcomes of these simulations and highlight the most relevant criticalities observed.

2.2.1 Pipe Rupture

In the following figures, the KPIs for all simulations are presented. Each simulation corresponds to the break of a specific pipe. Two distinct clusters can be clearly identified:

1. Cases where one or more demand nodes become isolated.
2. Cases where nodes remain connected but may still experience service degradation (and even complete service unavailability) due to low pressure.

Simulations were sorted by decreasing number of isolated demand nodes, and then by decreasing number of nodes affected by pressure reductions.

Figure 1 and Figure 2 show that **only 2,460 simulations (approximately 13% of all pipelines) resulted in at least one isolated demand node, with a maximum of 50 isolated nodes**. More than half of these cases involved just one node, 72% involved only two nodes, and 87% involved up to four nodes. **Only seven simulations (corresponding to a total pipe length of 884m) produced the most extensive isolation**. Figure 3 reports detailed data for this cluster of simulations, enabling – if needed – an additional level of prioritization that considers, for example, the potentially impacted demand. A further analysis step, focusing on all specific parameters and on the types and amounts of affected demand, can provide the full set of insights needed for informed decision-making.

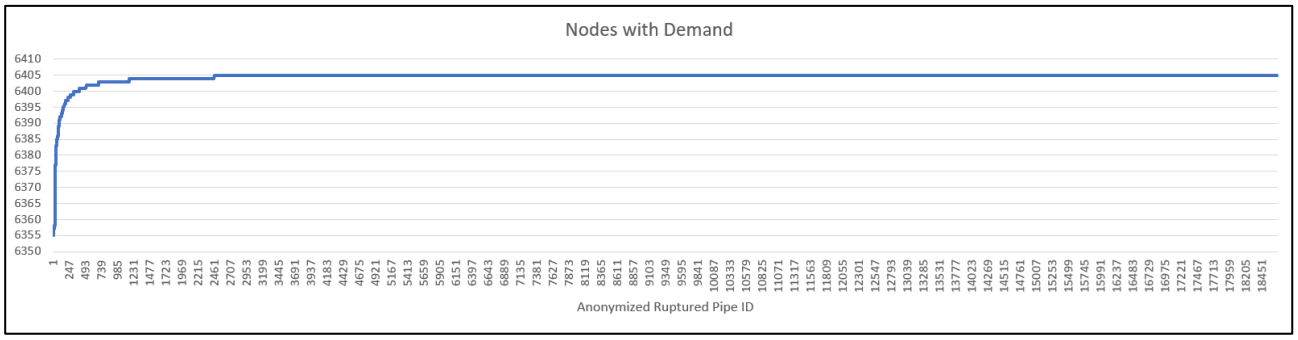

Figure 1: Number of nodes with demand per pipe-break simulation

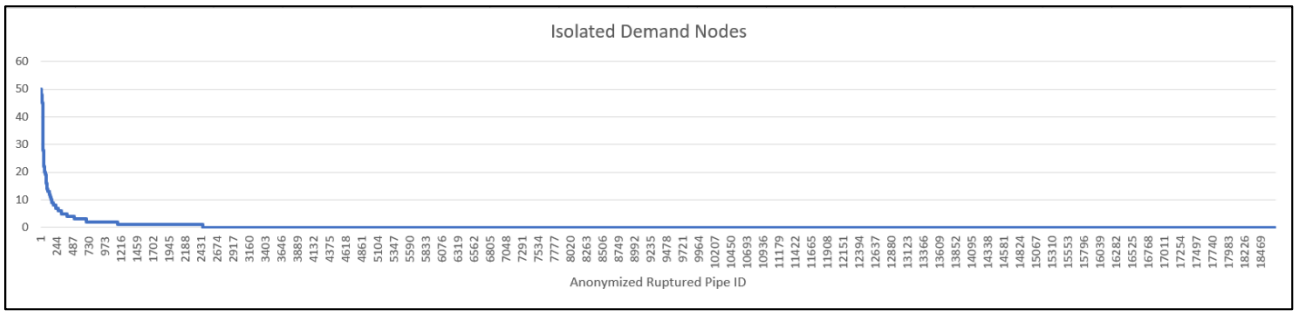

Figure 2: Number of isolated demand nodes per pipe-break simulation

Nodes with demand		Sim. PipeN-1		Pipes total	Unmet water demand (l/s)			Nodes with demand		Sim. PipeN-1		Pipes total	Unmet water demand (l/s)		
Connected	Isolated	N.	%	Length	MIN	MAX	Average	Connected	Isolated	N.	%	Length	MIN	MAX	Average
6355	50	7	0,04%	884,08	7,76	21,48	14,61	6386	19	7	0,04%	1002,56	1,83	14,17	6,31
6356	49	4	0,02%	268,52	3,37	7,85	5,49	6387	18	4	0,02%	591,41	1,72	10,85	5,90
6357	48	8	0,04%	506,54	3,12	3,44	3,33	6388	17	5	0,03%	98,65	0,00	10,28	7,27
6358	47	8	0,04%	485,81	3,19	3,22	3,19	6389	16	6	0,03%	877,65	0,00	9,44	2,64
6359	46	1	0,01%	10,37	1,81	1,81	1,81	6390	15	5	0,03%	1910,08	0,00	9,20	3,82
6360	45	4	0,02%	16,1	1,81	6,65	3,02	6391	14	11	0,06%	1293,72	0,00	13,66	6,53
6361	44	1	0,01%	433,87	6,65	6,65	6,65	6392	13	24	0,13%	2293,02	0,00	11,39	3,40
6362	43	1	0,01%	93,68	4,56	4,56	4,56	6393	12	12	0,06%	637,06	0,00	20,12	6,86
6367	38	2	0,01%	287,44	4,55	4,70	4,62	6394	11	16	0,09%	1399,61	0,00	21,66	7,66
6371	34	1	0,01%	255,02	4,34	4,34	4,34	6395	10	13	0,07%	94,93	0,00	16,26	8,17
6372	33	1	0,01%	8,41	4,07	4,07	4,07	6396	9	22	0,12%	3496,36	0,00	14,05	5,76
6377	28	4	0,02%	1053,36	2,98	4,17	3,77	6397	8	34	0,18%	2816,01	0,00	16,06	4,18
6379	26	1	0,01%	946,6	2,76	2,76	2,76	6398	7	39	0,21%	4202,78	0,00	17,17	5,48
6380	25	1	0,01%	43,33	2,26	2,26	2,26	6399	6	48	0,26%	4598,19	0,00	18,98	5,58
6381	24	5	0,03%	1529,7	2,18	7,40	3,24	6400	5	81	0,43%	6171,83	0,00	19,40	7,01
6382	23	1	0,01%	6,57	2,06	2,06	2,06	6401	4	108	0,58%	7108,39	0,00	21,94	7,23
6383	22	6	0,03%	1047,45	2,01	8,28	5,32	6402	3	186	1,00%	13171,2	0,00	23,19	6,37
6384	21	6	0,03%	1449,82	1,83	8,19	3,85	6403	2	469	2,51%	23975,92	0,00	121,54	7,79
6385	20	9	0,05%	975,2	2,08	14,24	7,95	6404	1	1299	6,95%	63261,82	0,00	115,92	7,07

Figure 3: Detailed statistics for the isolated-node cluster

In addition to node isolation which identifies the most severe connectivity issue, pressure conditions provide a second fundamental perspective on the system's response to pipe failures. Indeed, a pipe break also causes a pressure drop, which – when approaching zero – still indicates a complete loss of service. Figure 4 therefore reports the number of nodes with pressure $p \leq 0$ across the ordered simulations.

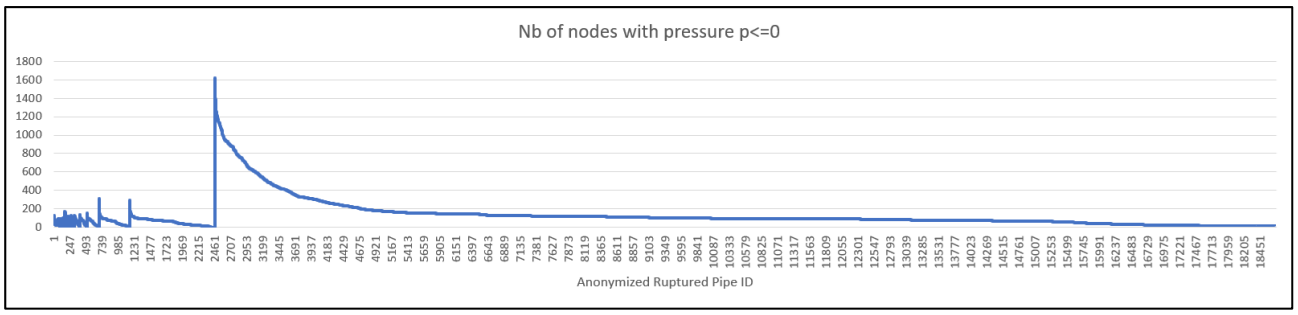

Figure 4: Number of nodes with p≤0 per pipe-break simulation

First, the isolated-node cluster (2.460 simulations resulting in the isolation of one or more nodes) clearly stands out. Although the initial simulations within this cluster correspond to the highest number of isolated nodes, and the final ones involve isolation of only a single node, two distinct peaks appear in the middle of the distribution. These peaks indicate a more widespread service disruption compared to the rest of the cluster and suggest the need for further investigation.

Overall, however, the extent of disruption within this cluster is significantly lower than in the roughly 6,000 subsequent simulations, which involve no isolated nodes. As expected, in almost all simulations a minimum subset of nodes (~100 nodes) experiences some degree of service degradation. Nevertheless**, a small, clearly identifiable group of pipes causes severe service loss ($p \leq 0$) to a very large number of nodes (up to 1.617).** Even in this case, the number of impacted nodes drops by half after only 330 simulations, and falls to 400 after an additional 780 simulations, covering a total of approximately 50 km of pipes – thus limiting the critical subset requiring deeper analysis for targeted mitigation.

The patterns observed for nodes with pressure below $p \leq 5$, $p \leq 15$, and, to some extent, $p \leq 30$, mirror those for $p \leq 0$, confirming consistent degradation as pressure decreases. Figure 5 shows the pressure levels across nodes for each threshold category, allowing the identification of reasonable optimisation targets depending on the emergency scenario. For example, in an emergency, maintaining a head of 15m may still be acceptable, even if lower than nominal operating values.

The good news is that 785 demand nodes remain fully operational in all simulations.

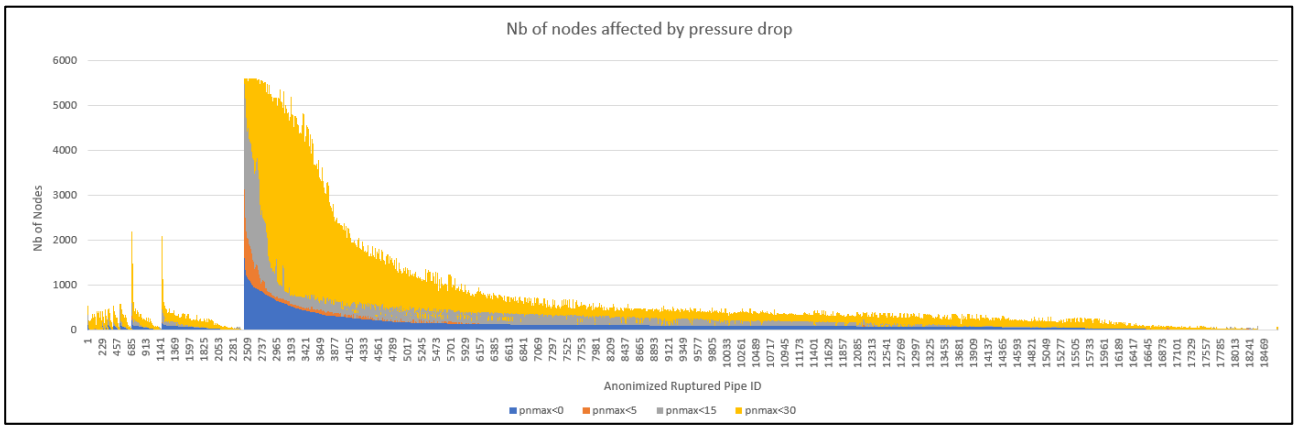

Figure 5: Pressure distribution throughout nodes per pipe-break simulation

When evaluating the unserved demand (see Figures 6 and 7), the most severe cases reach 87% of the total demand experiencing inadequate pressure. Fortunately, only 2.069 pipes cause disruptions affecting more than 20% of the total demand.

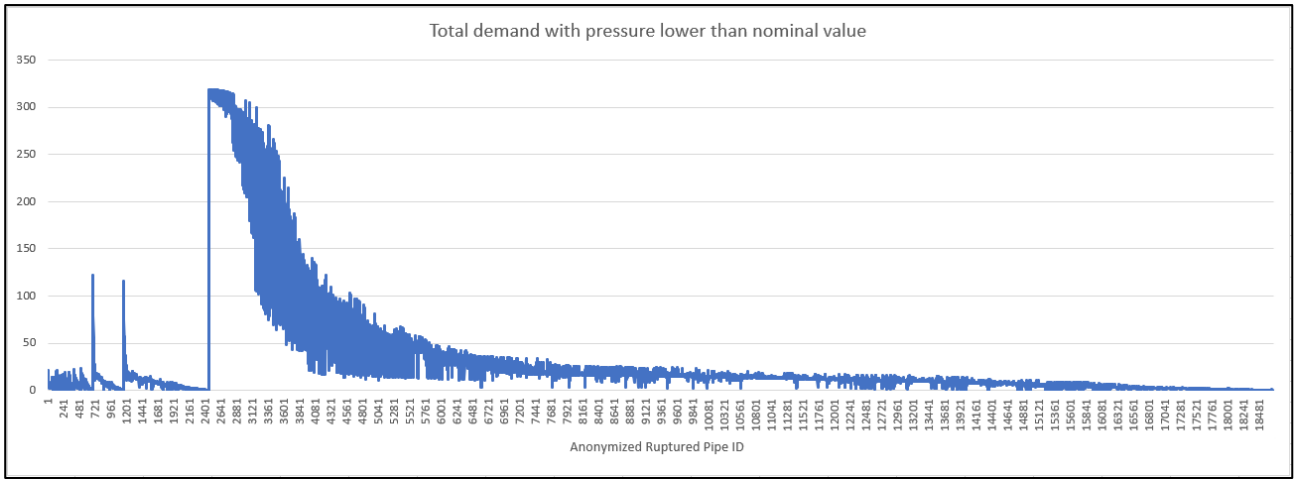

Figure 6: Distribution of disrupted demand severity across simulations

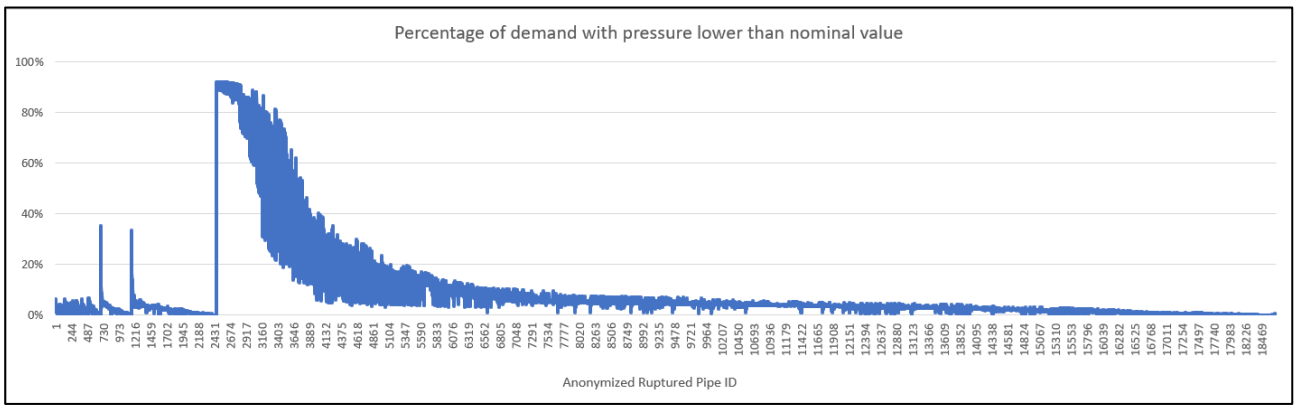

Figure 7: Percentage of demand affected by pressure drop

Regarding hydrants, the highest number of isolated hydrants (see Figure 8) corresponds – as expected – to the same cluster of simulations that produce isolated demand nodes. However, a non-negligible number of hydrants become isolated even in cases where overall network pressure remains adequate. This is due to the higher concentration of hydrants in certain areas, likely linked to the types of loads present.

When hydrants with inadequate pressure are also considered – a critical aspect for ensuring proper operation – the resulting trend (Figure 9) closely mirrors that observed for demand nodes. Figure 10, which reports the percentage of hydrants unavailable in case of need, further highlights that a limited subset of simulations can be particularly critical.

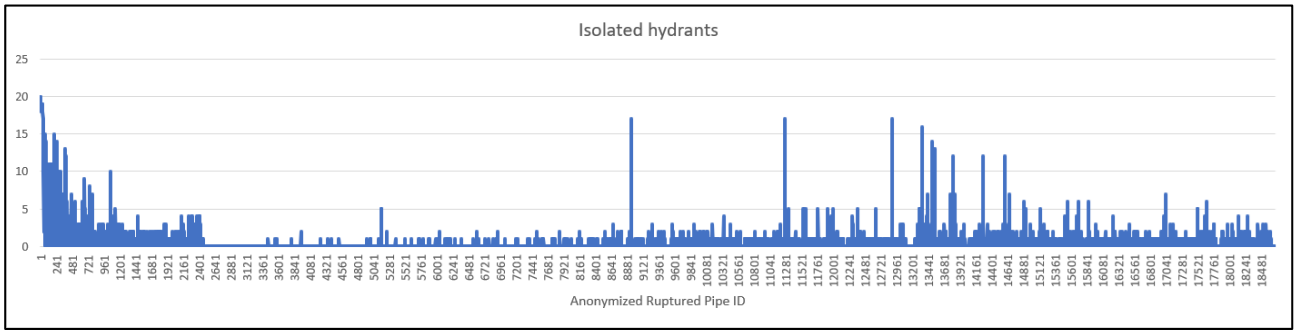

Figure 8: Number of isolated hydrants in pipe-break simulations

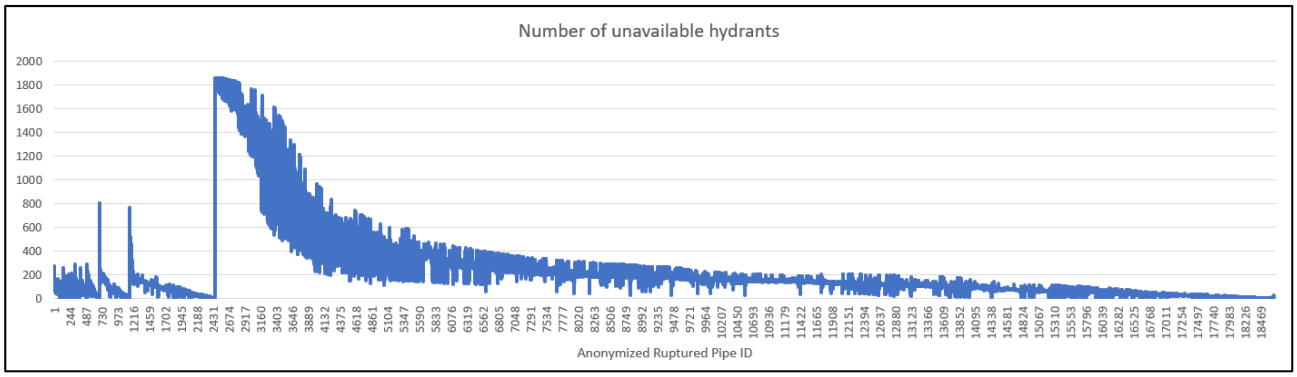

Figure 9: Number of hydrants with inadequate pressure

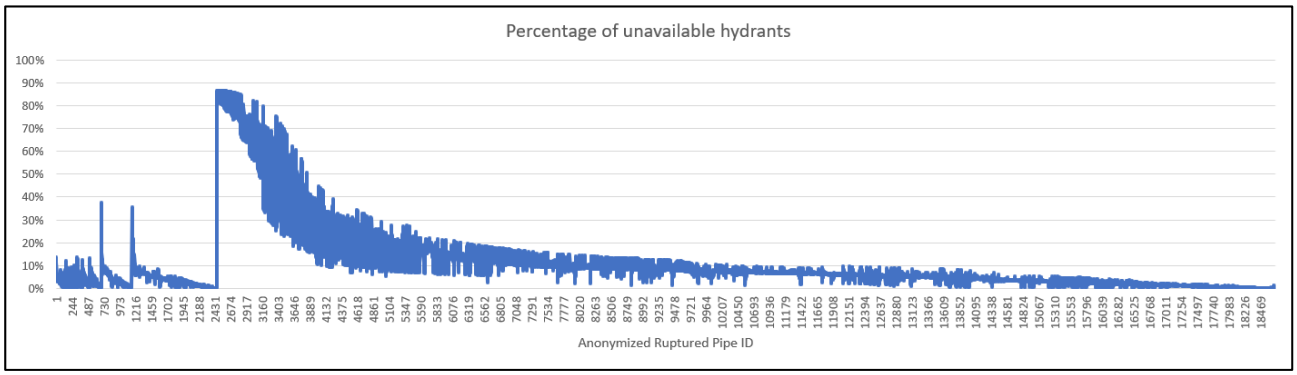

Figure 10: Percentage of unavailable hydrants

Another relevant parameter in pipe-break scenarios is breakflow, the volume of water leaking from a ruptured pipeline. Breakflow – shown in Figure 11 – behaves differently from the other KPIs: several singular points emerge, including a few critical cases in which normal pressure conditions coexist with extremely high

leakage. Such events may initially go unnoticed, yet they pose significant environmental risks and may trigger interdependency-related impacts on geographically co-located infrastructures.

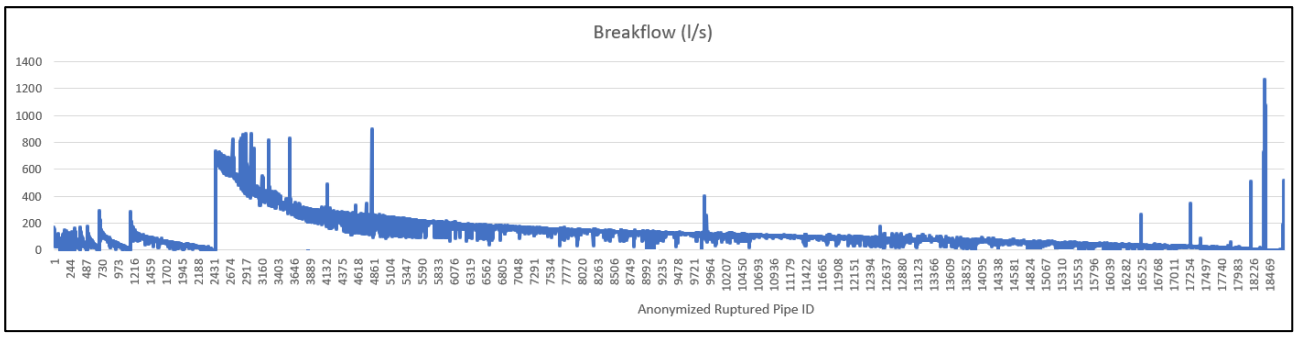

Figure 11: Breakflow per simulation

Finally, Figure 12 summarises the key values discussed above, enabling the identification of the number of cases that should be further investigated under different desired robustness levels (and corresponding intervention budget configurations).

Low pressure Nodes	Simul. PipeN-1	Unmet demand MIN	Unmet demand MAX	Unavailable hydrants MIN	Unavailable hydrants MAX	Breakflow MIN	Breakflow MAX
88%	40	92%	92%	86%	87%	610,98	736,78
87%	43	90%	92%	85%	87%	563,48	721,14
86%	55	89%	91%	82%	86%	513,64	729,27
85%	65	89%	91%	80%	85%	506,56	704,6
84%	31	88%	89%	79%	82%	434,75	692,25
83%	57	87%	88%	77%	82%	419,02	680,87
82%	47	86%	87%	76%	82%	489,22	675,61
81%	53	85%	87%	75%	80%	398,12	677,11
80%	28	82%	86%	74%	75%	441,73	716,43
79%	25	82%	85%	73%	74%	442,26	825,4
78%	36	80%	84%	72%	73%	421,42	540,87
77%	39	80%	83%	71%	72%	411,04	504,73
76%	36	79%	82%	70%	75%	365,15	475,93
75%	31	77%	81%	69%	75%	363,74	430,5
74%	25	77%	79%	68%	69%	380,7	477,48
73%	22	76%	79%	67%	70%	383,45	469,82
72%	23	75%	76%	67%	68%	375,15	813,58
71%	19	74%	77%	66%	72%	353,44	425,98
70%	25	73%	76%	66%	71%	350,41	832,96
69%	33	72%	74%	65%	66%	361,43	856,07
68%	39	71%	73%	64%	70%	346,55	827,5
67%	32	70%	73%	63%	65%	356,34	865,02
66%	16	69%	72%	62%	69%	342,97	637,2
65%	21	68%	70%	62%	67%	340,47	411,22
64%	17	67%	69%	61%	63%	347,53	426,49
63%	23	66%	67%	60%	62%	349,92	862,15
62%	19	65%	66%	59%	61%	344,92	389,75
61%	25	64%	65%	58%	60%	341,6	754,85
60%	22	63%	64%	57%	59%	338,93	519,2
59%	20	62%	65%	56%	62%	330,27	521,23
58%	15	57%	62%	55%	59%	332,27	509,21
57%	20	60%	62%	55%	61%	323,31	509,65
56%	11	59%	60%	54%	56%	328,28	494,16
55%	21	55%	59%	53%	55%	328,46	552,43
54%	16	57%	58%	53%	54%	329,16	358,97
53%	17	53%	57%	52%	54%	325,44	538,11
52%	20	51%	56%	51%	57%	313,32	477,75
51%	18	52%	55%	50%	57%	312,95	472,77
50%	19	50%	54%	50%	56%	311,01	460,79
49%	20	49%	52%	49%	55%	308,92	328,14
48%	15	48%	52%	48%	50%	314,87	458,36
47%	10	47%	50%	47%	48%	311,63	322,48
46%	10	46%	49%	46%	48%	309,92	454,87
45%	11	45%	48%	45%	46%	304,79	335,75
44%	12	44%	47%	45%	51%	301,94	323,58

Low pressure	Simul. PipeN-1	Unmet demand MIN	Unmet demand MAX	Unavailable hydrants MIN	Unavailable hydrants MAX	Breakflow MIN	Breakflow MAX
43%	12	42%	46%	44%	45%	306,32	365,33
42%	12	42%	45%	43%	45%	302,92	443,23
41%	14	41%	44%	42%	44%	301,56	435,02
40%	11	40%	43%	42%	43%	299,24	477,88
39%	18	39%	42%	41%	43%	296,37	469,98
38%	11	38%	39%	39%	42%	296,31	817,11
37%	25	37%	40%	39%	45%	288,73	419,63
36%	34	36%	39%	38%	44%	284,83	419,38
35%	29	35%	38%	37%	43%	283,34	412,86
34%	27	34%	36%	36%	39%	280,59	436,62
33%	31	33%	34%	35%	38%	284,05	431,46
32%	34	32%	35%	34%	39%	272,82	382,36
31%	36	30%	33%	32%	35%	266,64	372,31
30%	30	29%	32%	32%	35%	275,59	371,69
29%	54	27%	32%	30%	34%	255,92	365,88
28%	44	26%	32%	30%	35%	253,65	382,52
27%	64	25%	29%	29%	34%	0	349,38
26%	74	24%	30%	27%	33%	244,79	833,04
25%	63	24%	27%	26%	32%	244,93	345,32
24%	76	21%	26%	25%	31%	235,76	332,55
23%	55	20%	26%	24%	28%	243,8	491,5
22%	38	19%	23%	23%	29%	236,06	324,76
21%	75	17%	22%	22%	26%	219,49	340,28
20%	84	17%	22%	22%	28%	213,76	298,53
19%	79	16%	20%	21%	25%	207,8	293,54
18%	110	15%	18%	20%	25%	204,99	282,04
17%	116	14%	18%	19%	22%	199,24	280,25
16%	135	13%	17%	18%	23%	198,61	272,7
15%	170	12%	16%	17%	22%	191,42	274,29
14%	202	11%	16%	15%	22%	185,39	330,62
13%	184	10%	15%	15%	21%	176,6	346,87
12%	196	9%	14%	14%	20%	157,67	901,42
11%	259	8%	13%	12%	19%	144,74	261,01
10%	295	7%	12%	12%	18%	135,47	226,97
9%	263	6%	10%	11%	17%	131,09	213,79
8%	445	6%	9%	10%	16%	122,25	214,74
7%	883	5%	8%	8%	14%	100,63	198,54
6%	2019	4%	8%	7%	12%	81,2	169,37
5%	2103	3%	7%	5%	10%	64,36	399,73
4%	1927	2%	5%	4%	7%	44,1	118,81
3%	2047	1%	3%	3%	6%	29,95	96,71
2%	1565	1%	3%	2%	4%	0	72,14
1%	2259	0%	1%	1%	3%	0	346,22
0%	1351	0%	0%	0%	1%	0	1265,67

Figure 12: Overview of main KPIs

2.2.2 Pump failure

When assessing the impact of pump unavailability on the water distribution network, it is important to note that pump failures differ substantially from pipe breaks: no nodes become physically isolated, no breakflow occurs, and hydrants remain connected – although they may experience inadequate pressure.

Because the number of simulations required for this scenario is limited (the network includes only 54 pumping stations), the following cases were analysed:

1. Single-pump N-1 failure.
2. Total power blackout, which switches off all pumps simultaneously.
3. Single-pump availability for emergency restart.

Even when demand cannot be met, it remains constant and equal to the normal operating condition. Figure 13 illustrates this consistent behaviour across all three sub-cases.

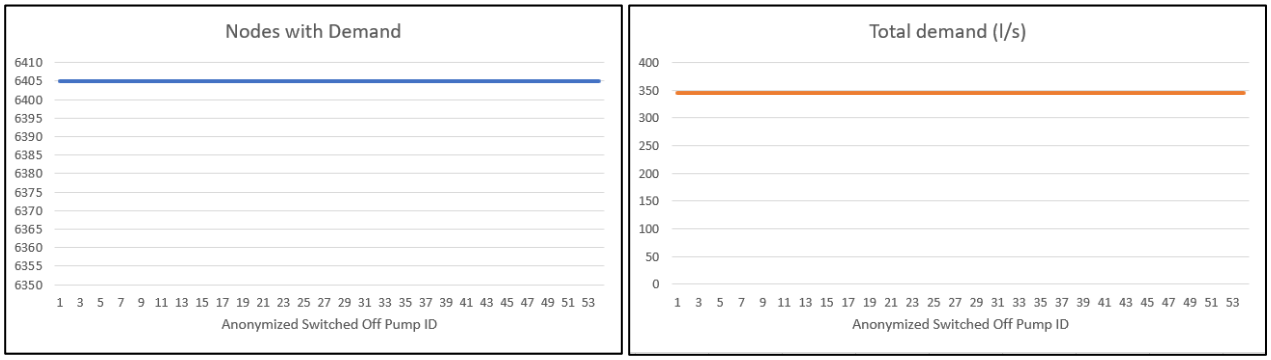

Figure 13: Demand nodes and Total demand across pump-failure scenarios

N-1 Pump

As already mentioned in the previous section, although nodes are not physically isolated, nodes experiencing very low pressure must be considered unserved. Figure 14 shows the number of unserved nodes in the case of pump shutdown. These are generally limited in number, indicating a predominantly local impact, which in some cases even disappears as it is compensated by the rest of the network. The overall pressure conditions across the network are shown in Figure 15, confirming the local nature of the disruption, which typically results only in a modest pressure reduction at some served nodes.

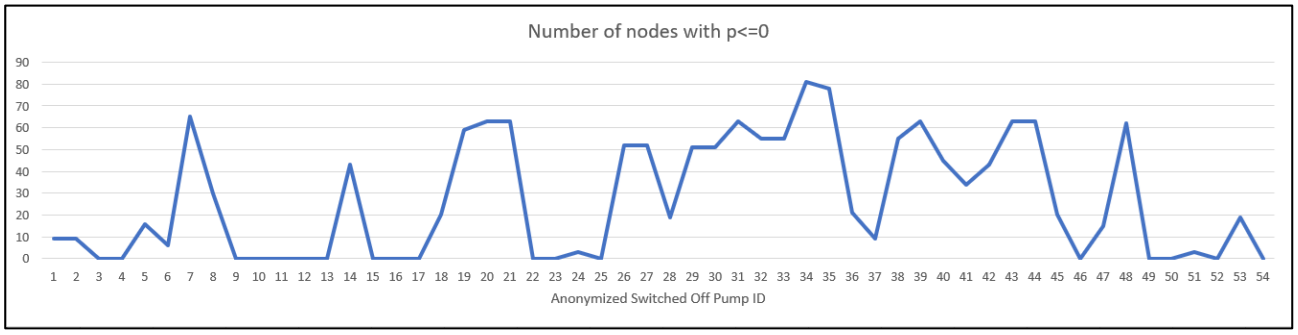

Figure 14: Not served demand nodes under single-pump failure

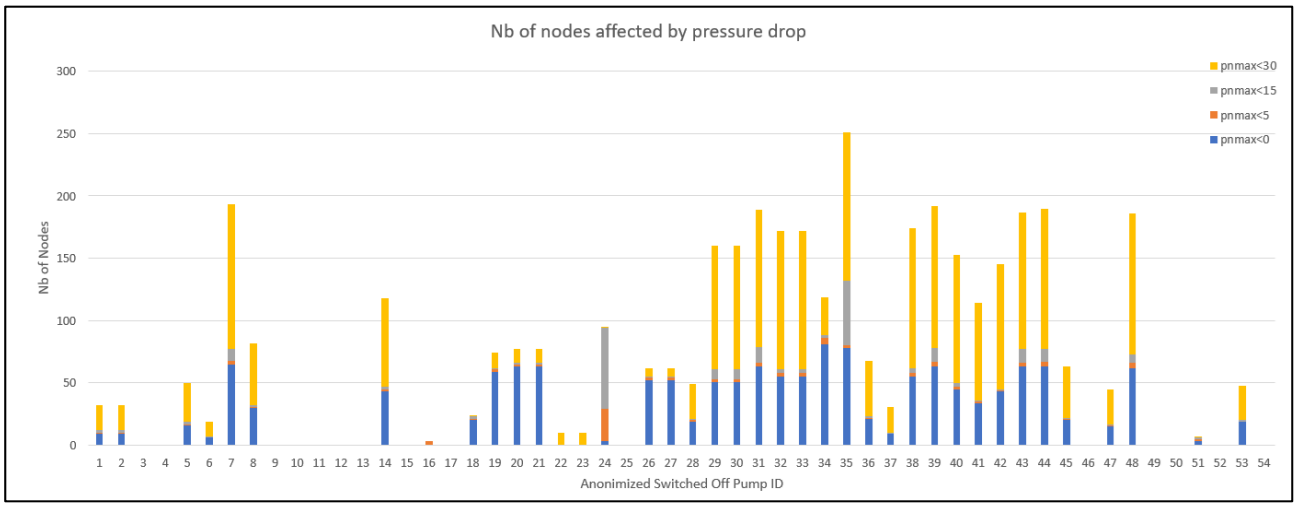

Figure 15: Pressure distribution across the network under single-pump failure

The overall impact on demand, shown in Figure 16, is modest, with the exception of a few cases associated with specific loads that may warrant mitigation.

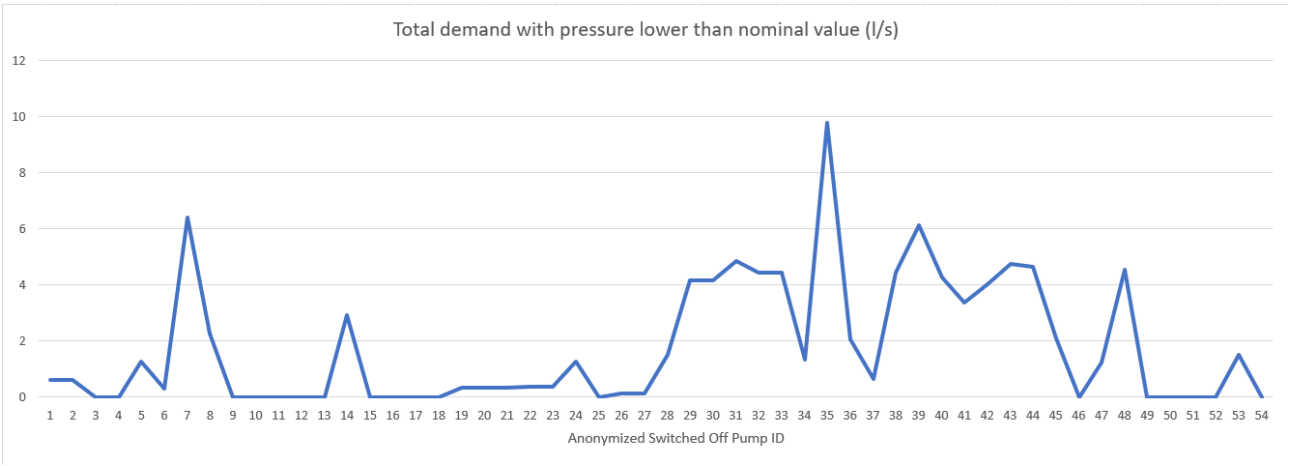

Figure 16: Demand with inadequate pressure under single-pump failure

The behaviour of hydrants – whose nominal pressure requirements are particularly critical – closely mirrors the trend observed for demand, as shown in Figure 17.

Taken together, the demand- and hydrant-related indicators reinforce the conclusion that **pump failures tend to produce only localized disruptions**.

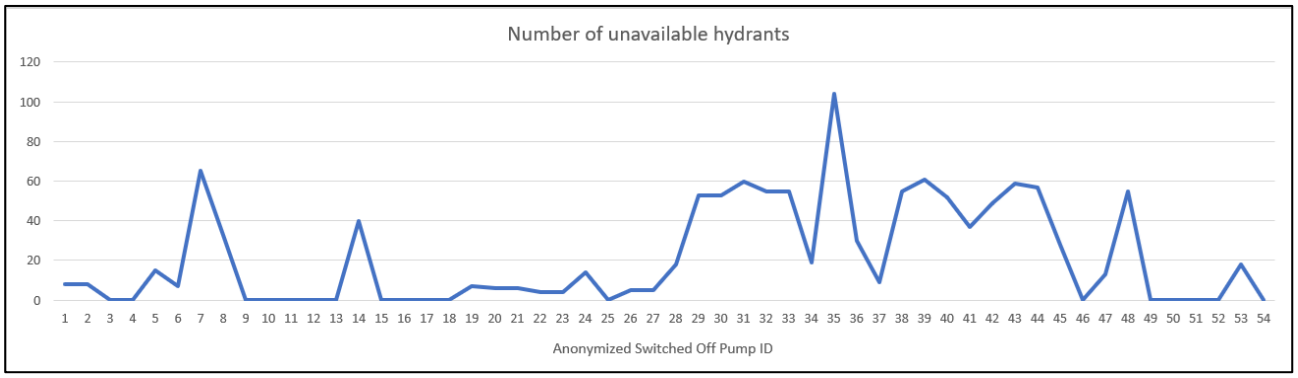

Figure 17: Hydrants with inadequate pressure under single-pump failure

Power Blackout

The power-blackout scenario drastically changes the situation. As shown in the summary provided in Table 1, **all nodes are heavily affected. Hydrants are also compromised**, making this a particularly dangerous condition in the event of explosions or fires triggered by natural hazards or NaTech accidents. It is important to note that, in real-world systems, storage tanks would at least partially and temporarily mitigate these effects but we cannot quantify their contribution as storage tanks are not included in this model.

Table 1 – Summary of key indicators under total blackout

	Nodes	Hydrants	Demand (l/s)
Isolated	0	0	0
p≤0	6330	2099	337,74
0<p≤5	75	39	7,66
<5p≤15	0	0	0
<15p≤30	0	0	0

Single Pump emergency operation

The scenario examined here concerns emergency management situations in which only a single pump can be re-powered using auxiliary generators. Utilities typically have 3–5 such generators available for emergencies, and optimisation studies are usually carried out to determine their optimal deployment, although this information was not available in this case.

As shown in Figures 18–19, supported by the presence of numerous wells, each pump is able to supply its own area (indeed, the number of nodes with pressure below the nominal value is always lower than the total of 6.405 demand nodes). Some pumps, however, can extend their influence more than others, although – when considering the overall pressure distribution across nodes – these differences remain limited.

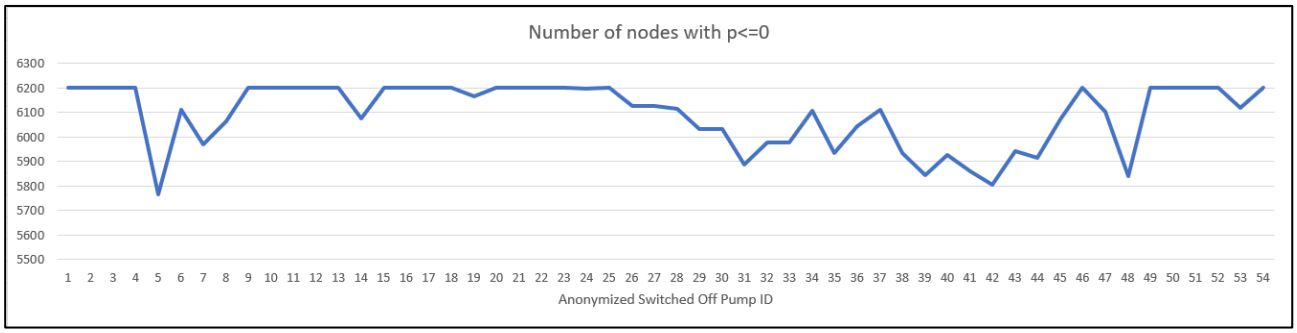

Figure 18: Demand nodes affected by service loss under single-pump emergency operation

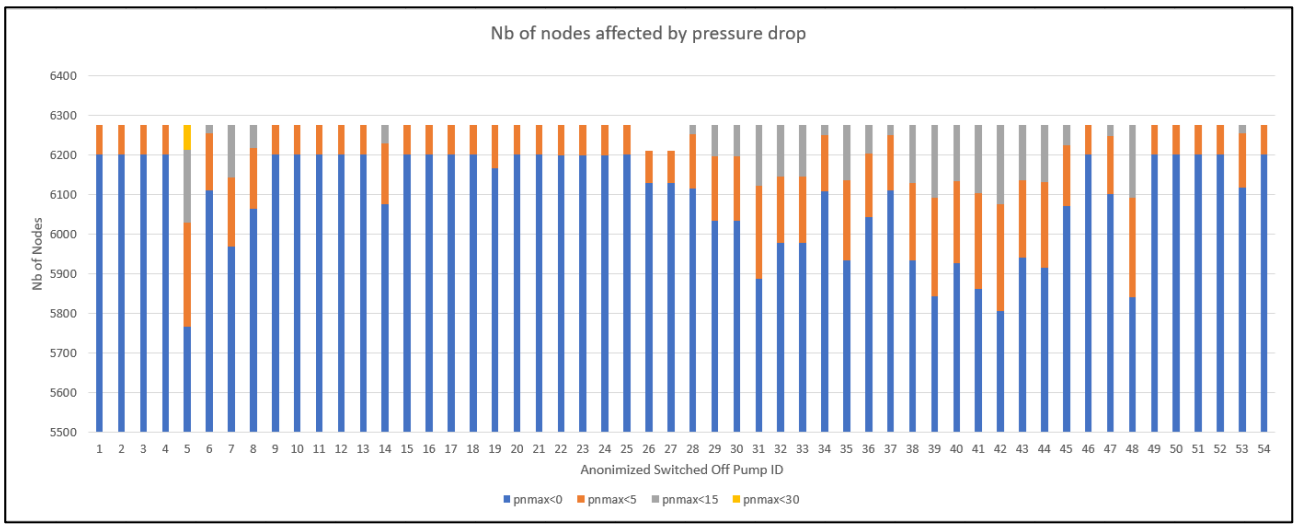

Figure 19: Pressure distribution in demand nodes under single-pump emergency operation

As expected, the served demand (see Figure 20) shows little variation among the different cases, and the same holds for hydrants (see Figure 21). **This suggests that prioritisation decisions may be better guided by the types of critical loads present in the area (e.g., hospitals) rather than by the number of users that can be resupplied**.

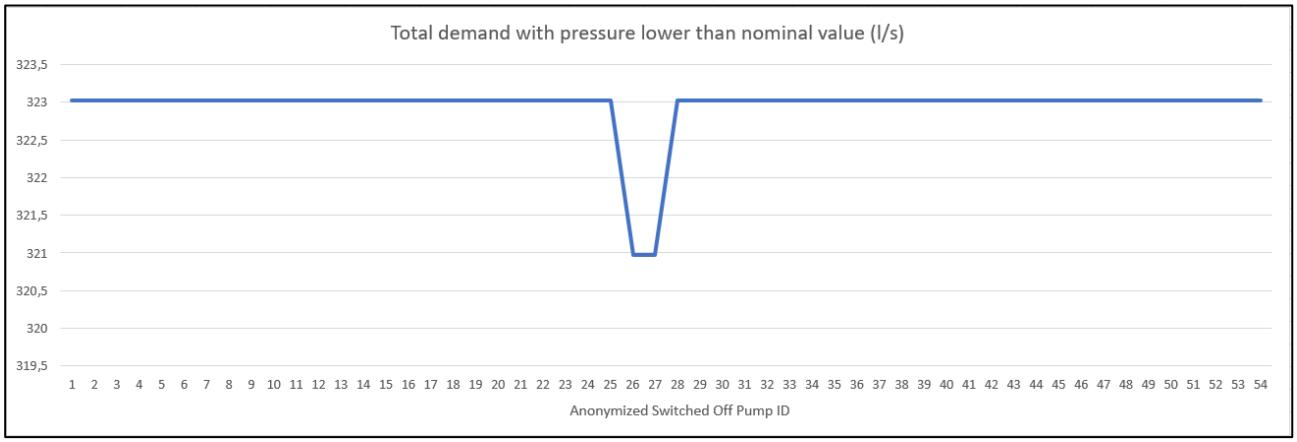

Figure 20: Unserved demand for single-pump restart

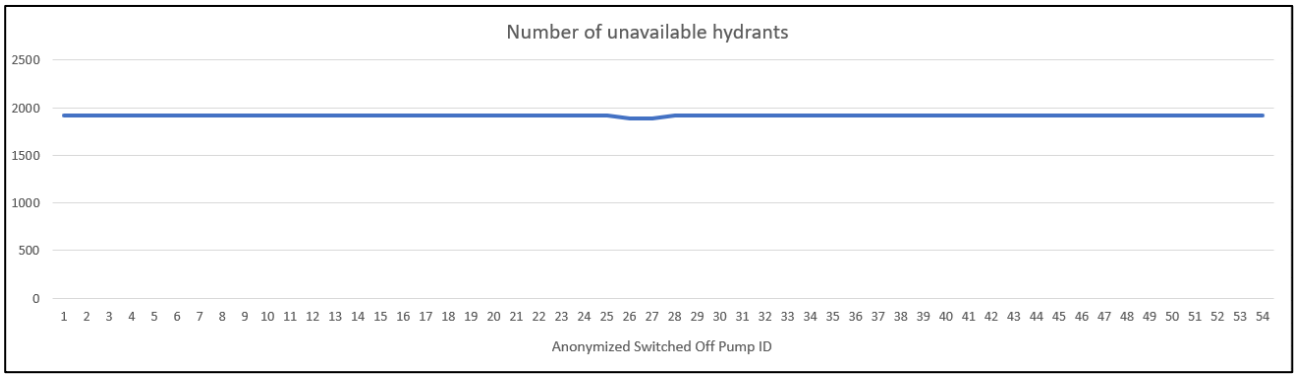

Figure 21: Unavailable hydrants for single-pump restart

2.3 Findings and Discussion

The insights presented above demonstrate that extensive, high-resolution simulations provide both a deep understanding of network behaviour and a risk-informed basis for prioritisation following the *As Low As Reasonably Practicable* (ALARP) principle – i.e., reducing risks to the lowest level reasonably achievable while considering costs, effort, and proportionality of mitigation measures.

Despite the size and complexity of the network, the simulation campaign shows that only about 10% of all possible failures are genuinely critical. This subset can be further reduced by prioritising the most severe cases and incorporating contextual considerations (e.g., presence of critical facilities). Although no network screenshots are provided for confidentiality, it is worth noting that many critical pipes are geographically contiguous, which simplifies maintenance and mitigation planning – while others would be extremely difficult to identify without simulation-based support.

Beyond node pressure, breakflow – often underestimated – deserves particular attention due to its potential cascading effects. Targeted replacement of vulnerable materials could reduce overall network vulnerability by up to 90%, a result of utmost relevance for natural-hazard scenarios such as earthquakes or liquefaction.

Dependence on the power grid emerges as another underestimated vulnerability. While no single pump is indispensable and localised power losses may have limited impact, a widespread blackout severely compromises the system, especially regarding hydrant availability during emergencies. Ensuring hydrant operability near critical loads is therefore essential. Furthermore, although not analysed in this study, telecommunication failures – which may accompany wide-area blackouts – could further hinder remote-control operations needed for emergency management and fault isolation.

Interpretation of the simulation results consistently benefits from local knowledge of the territory, which supports a more accurate assessment of criticalities.

Regardless of the specific numerical values, a central conclusion emerges: while operators often remark that no machine will ever outperform a skilled technician in many field situations, there are tasks – such as assessing 20,000 low-impact scenarios to confirm system robustness or identifying the 200 (or even 20) most critical ones – for which human capability cannot match systematic, exhaustive machine-based analysis.

Moreover, the simulation environment can support even more complex analyses, such as evaluating all possible combinations of 30 simultaneous pipe failures among 200 pipelines in a liquefaction-prone area - an activity initiated within the ReturnLand/ReturnVille virtual test bed.

The next step necessarily requires human expertise: once a shortlist of critical cases has been identified, expert judgement becomes essential, with simulation acting as a support and validation tool rather than a replacement for decision-making.

Mitigating this small set of critical cases – often through relatively minor investments – can dramatically enhance service reliability and emergency preparedness. As a general recommendation, attention should be directed toward large loads, path redundancy, and components so robust that a failure may initially go unnoticed.

Finally, all simulation results can be linked to the GIS representation of the network and made accessible to operators through external decision-support systems, without requiring direct access to Infoworks WS PRO. In this sense, once the simulation campaign has been completed, the resulting GIS-based representation can be reused and queried independently of the simulation environment. With regard to the effort required to build such a dataset, the simulation campaign needed several months to reach a stable configuration; however, once the workflow is consolidated, running the simulations required to populate the GIS representation would in principle take 30–40 days. Further reductions in processing time would be possible by increasing the number of available software licences, although these would typically be needed only for a limited period.

For this reason, should this approach become widely adopted, it may be advantageous to equip a centralised computing platform (e.g., a High-Performance Computing infrastructure) with a pool of licences and offer the simulation capability as a service, rather than replicating short-term resources at individual operator level.

3. Railways

This section presents the analysis of the simulation results for the railway network.

Section 3.1 provides a brief recap of the main methodological aspects – already discussed in Deliverable DV6.2.3 – for the reader's convenience, while also introducing additional modelling details that were not available at the time DV6.2.3 was released due to delays in the provision of some datasets. These additions are relevant to illustrate the full potential of the simulation environment and are required for the real case study. Section 3.2 describes the simulations performed and discusses the results, while Section 3.3 outlines the key findings and suggests future developments.

Once again, as this deliverable is public and the information may be sensitive, the discussion is deliberately kept general.

3.1 Methodology

The simulated network, described in Deliverable DV6.2.3, is the in corridor of the Tyrrhenian Calabria coastal stretch from Naples to Reggio Calabria and consists of:

- 85 medium-sized stations;
- 9 large stations;
- 750 km of double-track railway;
- 160 km of single-track railway.

The railway network model used for the simulations was developed based on precise and detailed information – protected by Non-Disclosure Agreement (NDA) – concerning train services and their main characteristics, as well as stations and their layouts, signalling systems, speed limits, gradients, and route lengths. This approach made it possible to fully exploit the computational power and accuracy of the selected simulator, OpenTrack, which employs differential equations to compute train speeds and positions and can be remotely controlled via its API, as discussed in Deliverable DV6.2.3. The implemented timetable represents a typical winter weekday.

To evaluate the impact of different hazards on the infrastructure, the effects of each hazard on the network must first be translated into corresponding anomaly events (or faults). In the railway domain (see Deliverable DV6.2.1), an anomaly event is defined as an occurrence that results in at least one of the following situations:

- Unavailability of sections of line or plant (which reduce normal capacity both compared to circulation of trains with respect to plant services), accounting for tracks, traction or signaling.
- Capacity reduction (also considering the limitations in the performance of railway vehicles).
- Degradation of the functioning of the infrastructure (including sudden slowdowns, slowdowns beyond the rate and code reductions beyond the rate).

Next, the identified event or fault must be translated into the software framework. In OpenTrack, this may involve modifications to trains, track sections, or signalling systems. The faults modelled in OpenTrack may correspond either to infrastructure-related failures, such as track damage or signal malfunctions that interrupt train operations, or to operational disruptions, such as imposed speed restrictions or unplanned train delays.

Finally, it is necessary to define how the available software tools can be effectively used to manipulate the infrastructure and simulate these failures. Table 2 summarises the hazards considered and their possible translation into OpenTrack's modelling framework.

Table 2 – Modelling hazards in OpenTrack

Hazard	Effect on infrastructure	OpenTrack modelling
Wind	Speed reduction	Speed reduction on impacted edges
	Track blockage, due to tree collapse	Edge(s) blockage
	Power outage, in case a powerline is damaged, by wind itself or tree collapse.	Signal off with speed reduction up to 0
		Traction failure with edge blockage
River flooding	Bridge inaccessibility	Edge blockage
Soil liquefaction	Track/bridge deformation	Edge blockage
Landslide	Track blockage	Edges blockage
	Power outage	Signal off with speed reduction up to 0
		Traction failure leading to edge blockage
	Torn-off signal	Signal off
Wildfire	Speed reduction	Speed reduction on impacted edge(s)
	Track damage	Edge(s) blockage
	Signal damage	Signal not available with speed reduction up to 0
	Power outage (or disconnection)	Signal off with speed reduction up to 0
		Traction failure leading to edge blockage
Earthquake	Traffic halted (locations with PGA > 0.1g)	Edge(s) blockage
	Traffic halted (for inspection)	Edge(s) blockage
	Track damage	Edge(s) blockage
Tsunami	Traffic halted	Edge(s) blockage
	Track damage	Edge(s) blockage
Coastal erosion	Track damage	Edge(s) blockage

In summary, most of the hazards listed above result either in the interruption of a single edge or of groups of edges, or in a reduction of train speed along specific sections of the infrastructure. The analyses presented

in this deliverable focus on assessing the impact of the infrastructure in terms of single-edge interruptions only. By contrast, within the Proof of Concept Calabria, more complex scenarios were analysed, involving the simultaneous blockage of multiple edges – even at geographically distant locations – in order to assess higher-level, system-wide impacts, such as the closure of entire railway sections.

Simulations involving speed reductions on specific track sections could not be performed within the available timeframe, mainly due to delays associated with the NDA process. In addition, information on speed limits applied to trains in the event of malfunctions is highly dependent on the nature of each disruption. In the case of signal malfunctions, OpenTrack currently does not allow signal states to be modified via the API. As a result, the only viable workarounds consist in implementing the railway operator's operational policies – either by stopping trains through edge blockage, imposing reduced speeds, or simulating manual traffic dispatching by a human operator.

It should also be noted that OpenTrack does not allow track edges to be blocked directly via the API, but only through the graphical user interface (GUI). Through the API, it is nevertheless possible to enable or disable edges that have been pre-defined and blocked in the GUI. This approach is highly time-consuming, especially considering that the model includes approximately 5.400 edges.

To overcome these limitations, a novel programmatic workaround was developed. Starting from an OpenTrack-generated file listing all track edges, a graph-based representation of the railway network was constructed. This graph enables systematic navigation of the infrastructure and supports the identification, for each edge, of the corresponding routes, defined as ordered sequences of vertices representing valid train paths. Leveraging this structure, a specific OpenTrack API function – allowing the assignment of speed limits to selected sections of a route – was exploited. By imposing a speed limit equal to zero, this mechanism effectively replicates the functional effect of blocking a track segment, without requiring manual intervention through the graphical user interface. This approach enables full automation of failure injection, even in large-scale models comprising thousands of edges.

In the simulations presented here, the imposed speed restrictions were maintained for the entire simulation duration in order to analyse network behaviour under severe disruption conditions. Beyond the specific case study, this approach provides a general and reusable method for simulating a wide range of infrastructure failures in OpenTrack, enabling scalable, reproducible, and systematic analysis of railway network vulnerabilities.

It should be noted that it was possible to reliably identify the routes associated with most, but not all, edges. This limitation is inherent to the network structure: in several sections, edges lie on parallel tracks with identical lengths and the same number of segments, making it impossible to unambiguously determine which parallel edge is actually traversed. Such distinctions can only be resolved through the graphical user interface (GUI), which could not be used for this purpose within the timeframe of this deliverable.

Moreover, due to the computational time required for each simulation, the total number of simulations was reduced by limiting the analysis to edges located on branching tracks. This assumption does not significantly affect the validity of the results, as blocking edges sequentially along a single track is expected to produce equivalent outcomes in most cases.

Among the outputs of the simulation campaign are train arrival and departure times at each station, as well as deviations from the planned timetable (provided as an input to the model). This enables the assessment of delays and service interruptions for each train and the estimation of the number of affected passengers, in accordance with the KPIs defined in the previous deliverables.

To meet the objectives of the project, several OpenTrack outputs required post-processing. Owing to its planning-oriented design, OpenTrack primarily focuses on differences between scheduled and actual arrival and departure times of train services at stations, expressing these deviations in seconds. However, this approach overlooks some of the most relevant situations for the present analysis – namely, courses that do not reach their final destination. For such cases, OpenTrack reports a deviation of zero seconds between scheduled and actual times, which is misleading from a resilience and service-continuity perspective.

During the post-processing phase, these services were therefore explicitly identified and classified as follows:

- On-time;

- Early;
- Late;
- Not arrived;
- Not departed;
- Not stopping in the network segment.

These enhancements were considered both necessary and appropriate, as they not only improve the quantification of the severity of disruptions affecting the infrastructure but also enable a clearer assessment of their nature. This, in turn, supports the definition of more effective and targeted mitigation strategies.

3.2 Simulations and Results

3.2.1 Simulations overview

The simulations were performed under the following parameter settings:

- Simulation step: 1 sec;
- Simulation rate: Best;
- Delay scenario: none;
- Adhesion: good;
- Mean delay: 0 sec;
- Performance: 100%;
- Start time: 0 seconds, corresponding to 00:00:00;
- End time: 86399 seconds, corresponding to 23:59:59.

The simulation step was set to the lowest possible value, and the simulation rate to the recommended value, in order to achieve higher numerical precision. This parameter defines the frequency at which each train computes its position, speed, and acceleration.

The start time and end time specify the beginning and end of the simulation (in seconds). Each simulation covered a 24-hour horizon, ensuring that all train services operating on the infrastructure were captured.

The delay scenario and mean delay parameters represent the average initial delay applied to trains (expressed in seconds).

The performance parameter, which scales train acceleration and maximum speed, was set to its maximum value in order to simulate trains operating under optimal conditions.

In total, 1139 simulations were completed and analysed, with an average duration of 31 minutes and a modal value of 22.2 minutes. Figure 22 reports the simulation duration (in minutes) for each processed edge.

The computation time required by OpenTrack for each simulation depends on several factors. The first is the simulation step: for instance, increasing it to 5 seconds significantly reduces the overall computation time, at the expense of temporal resolution, as some events or commands may no longer be captured.

A second factor is the duration of the simulated time window, which – when using the API – is limited to a maximum of 24 hours. In principle, computation time could be reduced by initially running shorter simulations (e.g., covering only peak or mid-day periods) and then re-running the most critical scenarios over the full 24-hour horizon. However, this strategy was not adopted due to a limitation of OpenTrack: trains whose scheduled departure time falls outside the simulated time window are excluded from the

simulation. As a result, restricting the time window would prevent all scheduled trains from being considered.

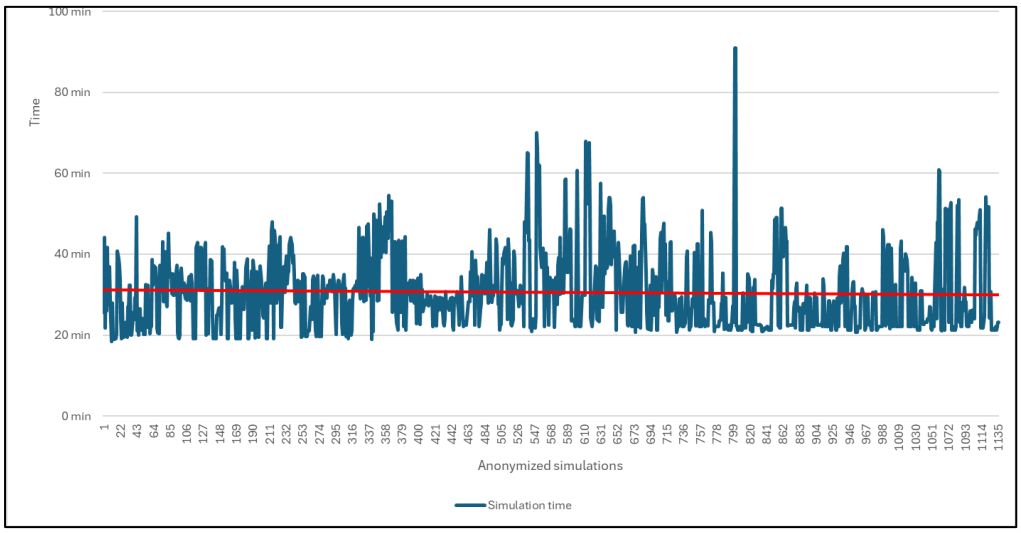

Figure 22: Simulation time

Another factor influencing computation time – clearly visible in Figure 23 – is the operational importance of the blocked edge. The greater the number of trains delayed or stopped as a consequence of a blockage, the longer the simulation tends to take to complete. Finally, computation time is also affected by the intrinsic performance of OpenTrack and the hardware on which it is executed, as well as by potential memory saturation when a large number of simulations are run consecutively.

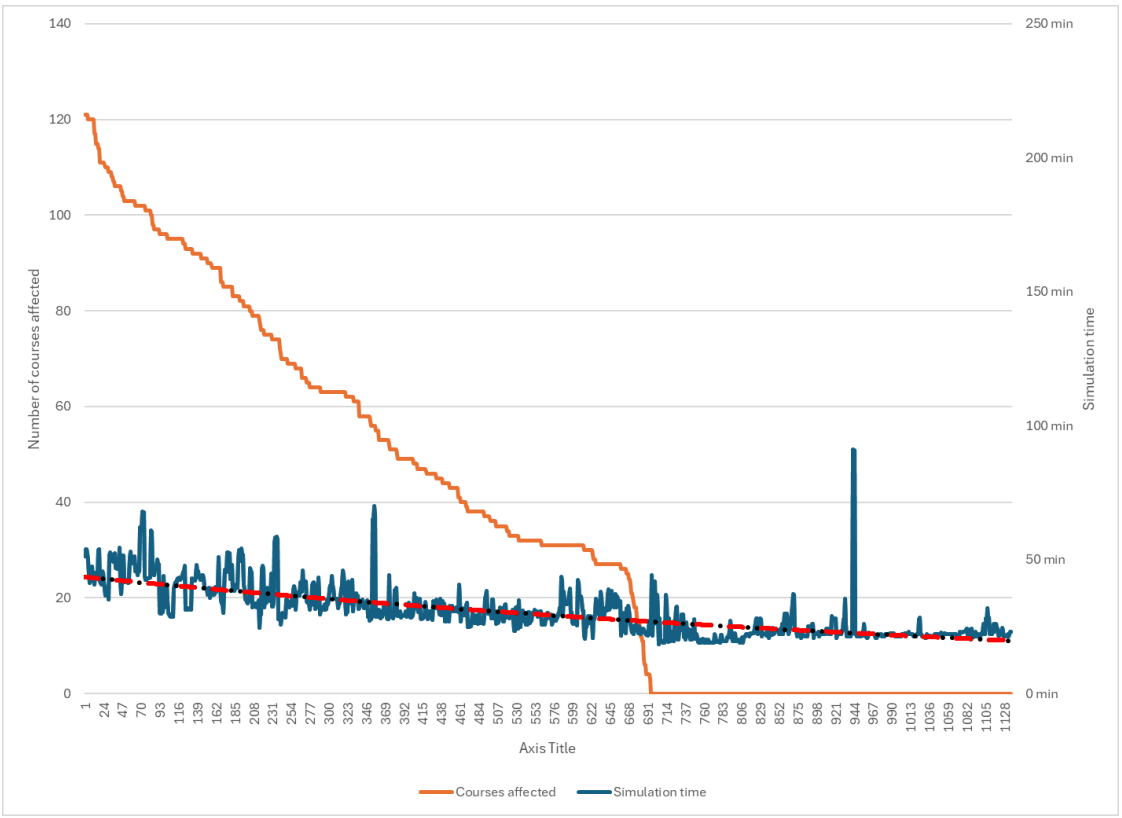

Figure 23: Simulation time and affected courses

3.2.2 Results

The following charts (Figures 24 and 25) present two complementary views of the relationship between the number of train services affected by the blockage of a specific edge and the number of services that would normally traverse that edge in the reference (i.e. no-failure) simulation. The two figures differ only in the ordering criterion, which is intended to better highlight the relative role and magnitude of the two aspects.

As expected, when an edge is blocked but no services traverse it under normal conditions, the number of affected services is zero. More generally, a clear correlation emerges: the higher the number of services passing through an edge in normal operation, the larger the number of services affected by its blockage.

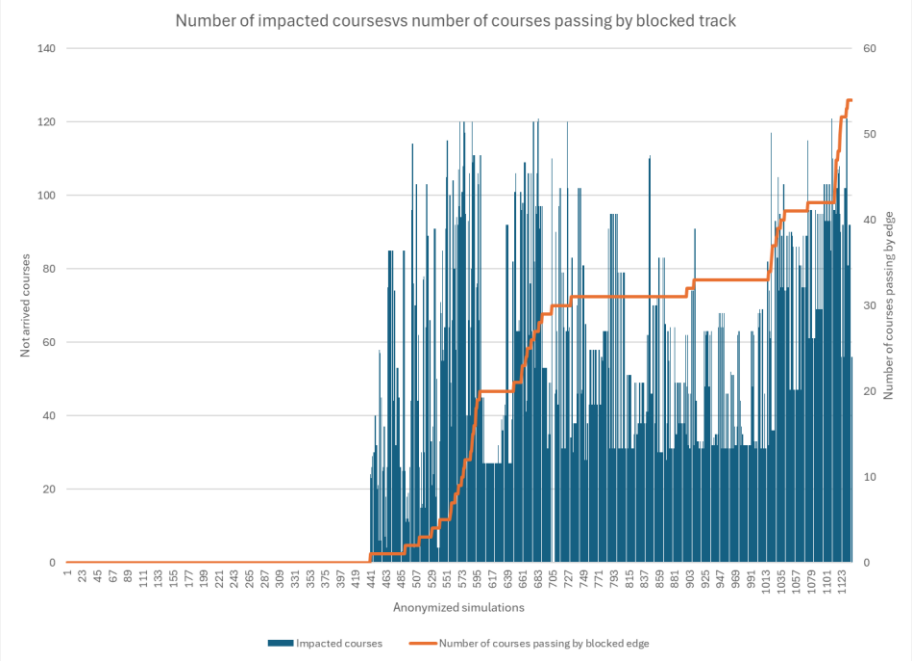

Figure 24: Number of courses impacted and courses passing by blocked track

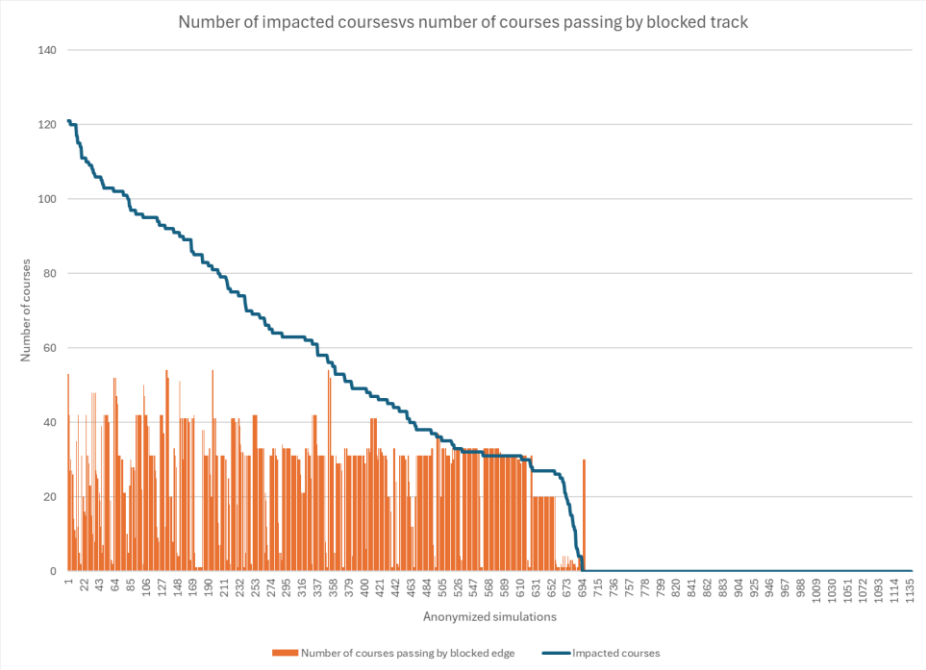

Figure 25: Number of courses impacted and courses passing by blocked track

Beyond the general behaviour observed, several cases deserve closer examination.

For example, simulations around index 694 in Figure 25 show that, although many courses normally traverse the affected edge (30 courses), none are actually impacted. These edges are located within or near station areas and have parallel tracks that can be used in case of congestion. It should be noted that, in this initial phase, only a limited number of alternative routes were implemented due to time constraints and because operational re-routing decisions – while based on established procedures – are carried out in real time by personnel in charge, incorporating ad hoc assessments of multiple factors (e.g., delays, congestion levels, and the nature of the disruption). Accordingly, this phase focused on assessing the direct impact of a blocked edge and on identifying potential mitigation strategies, rather than on fully modelling dynamic re-routing.

The first group of simulations in Figure 24 also highlights edges with very few normally traversing courses (2) that nevertheless result in a high number of impacted services (114). These cases are not anomalies: they correspond to transversal edges connecting parallel tracks in the vicinity of medium- or large-sized stations. Even minor congestion on such edges can propagate and generate widespread delays. Moreover, adjacent edges – on which a much larger number of services normally operate – produce a similar number of impacted services. This reinforces the interpretation that, in these situations, it is the broader station area, rather than an individual edge, that requires particular attention.

As shown in the previous charts, out of the 1.139 simulations, approximately 440 result in no impacted courses. For the remaining simulations, the number of impacted courses varies depending on both the position of the edge and the number of courses traversing it under normal conditions.

As mentioned in the previous section, in order to obtain results as accurate as possible – and also in consideration of project deadlines – the simulations were performed only on the railway section between Salerno and Reggio Calabria. Even when considering this reduced section and an average simulation duration of 31 minutes, running 1139 simulations continuously (24/7) would require at least 24 days. This estimate does not account for the interruptions that typically occur over such extended execution periods, whether due to manual intervention (e.g., development or debugging activities) or to performance degradation caused by intensive resource usage. In the specific case of this project, completing the full simulation campaign required approximately seven weeks, once all implementation issues had been resolved.

It should be emphasised, however, that the methodology described in the previous section is independent of the specific railway section analysed and is therefore applicable to any railway network, regardless of size or geographical location. The graph construction process relies on files generated by OpenTrack and adapts automatically to the implemented network model. The same applies to the simulation workflow and to the considerations discussed below.

For the analysed section (Salerno–Reggio Calabria), a total of 233 train services are included, belonging to the following five categories:

- High-speed trains (Eurostar);
- Long-distance trains (Intercity);
- Regional trains (Regionale);
- Freight trains (Merci);
- Deadhead trains (Invio).

In the reference simulation, twelve services – belonging to the high-speed, regional, and freight categories – originate at a station within the network but do not stop at any other station along the analysed section, only passing through it. These services are therefore not included in Figure 26, but are clearly visible as purple bars in the corresponding detailed charts (Figures 27, 28, and 30).

Once again, the results in Figure 26 show that, from approximately simulation 694 onward, no services are impacted. Conversely, in the initial simulations, the number of impacted services—considering both delays

and interruptions – reaches 100 out of 221 services (i.e., excluding the twelve non-stopping services discussed above).

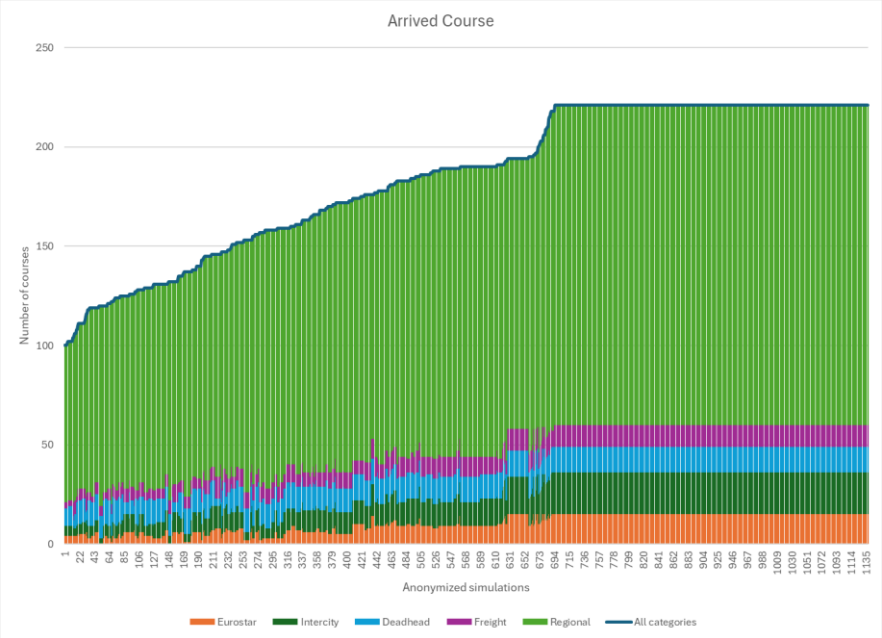

Figure 26: Total arrived courses

Figures 27 to 31 present detailed charts of the impacted services, disaggregated by train category.

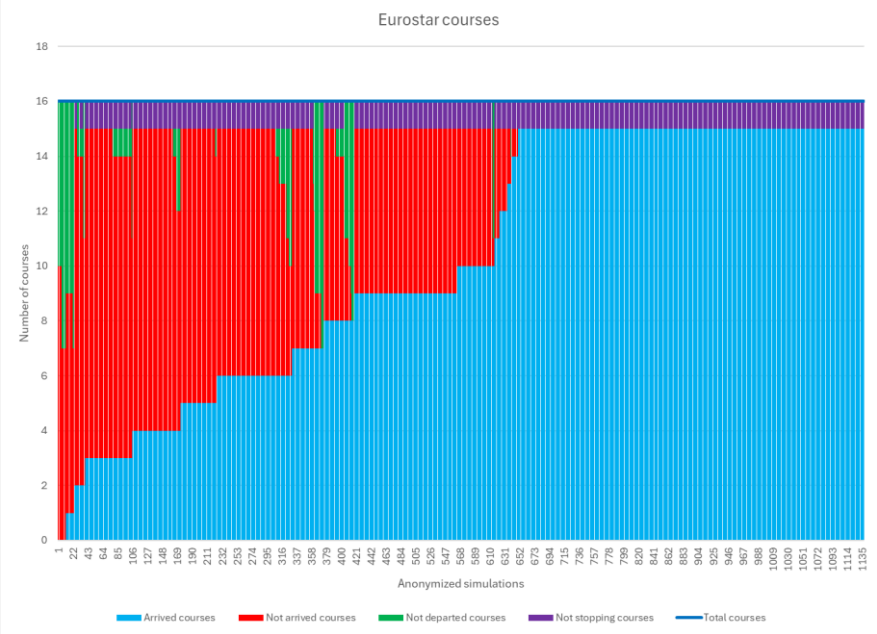

Figure 27: Overview on high-speed courses

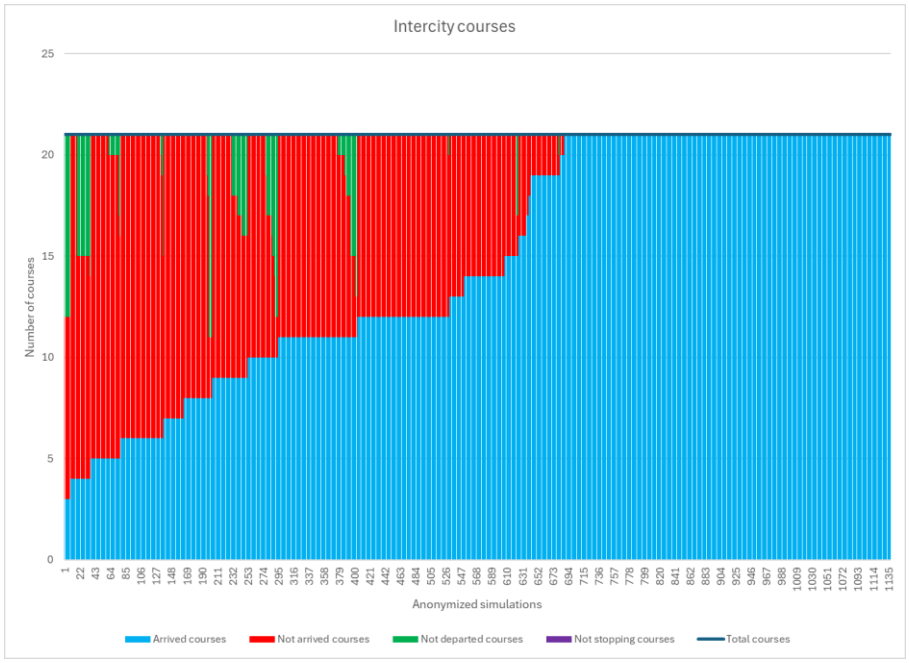

Figure 28: Overview on long-distance courses

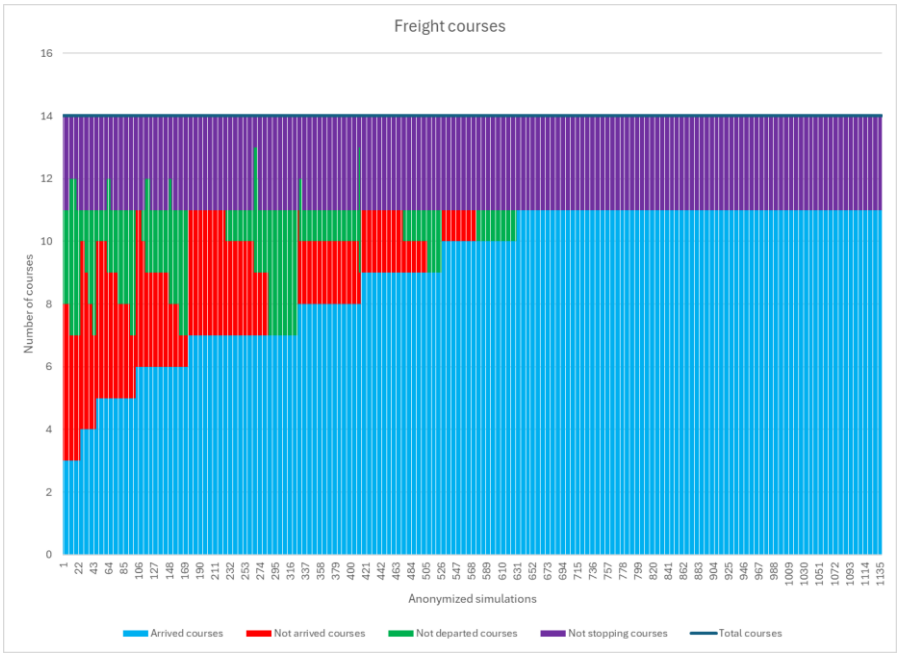

Figure 29: Overview on freight courses

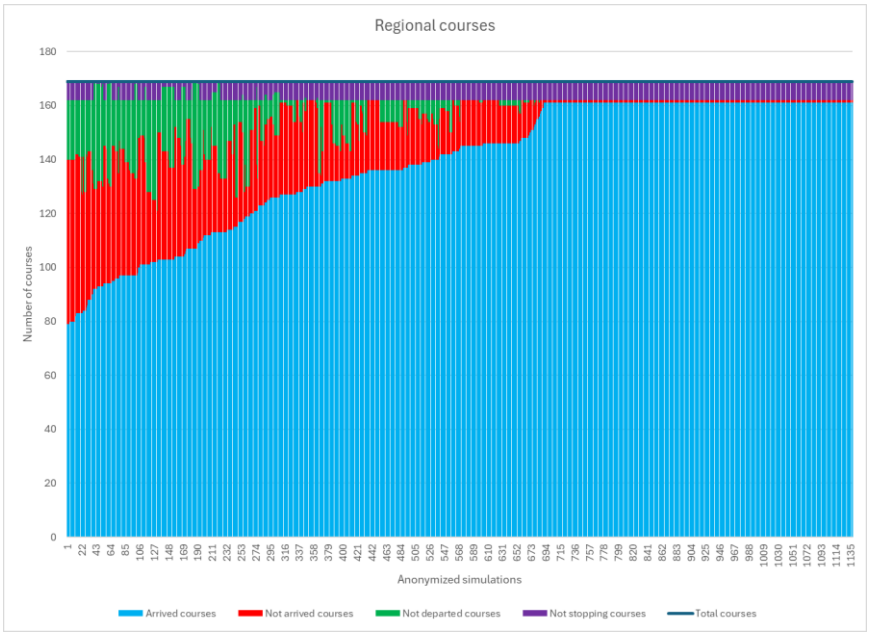

Figure 30: Overview on regional courses

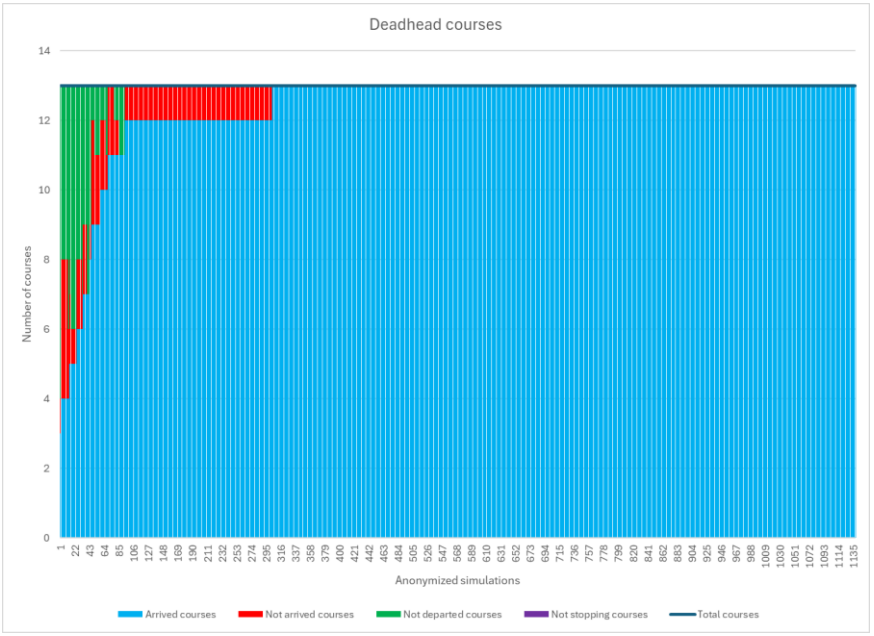

Figure 31: Overview on deadhead courses

The red bars, labelled *not arrived courses*, represent the number of courses that did not reach their final destination during the simulation timeframe. This may occur either because trains were forced to stop due to severe congestion and would not resume operations until the damaged railway section (i.e. the blocked edge) was fully restored, or because they experienced substantial delays and failed to reach their destination within the 24-hour simulation window – although they would have eventually arrived on the following day with significant delay.

The green bars indicate the courses that, due to congestion and/or blocked railway sections, were unable to depart from their origin station or, once departed, failed to reach the first scheduled stop.

The light-blue bars represent the courses that successfully reached their final destination - either on time or with delays – within the 24-hour simulation period.

Inspection of the previous charts shows that, particularly for high-speed and long-distance services, a substantial number of trains are affected, both in terms of courses that do not reach their final destination and courses that fail to depart from their origin station.

A more detailed examination reveals that many of these cases involve blocked edges located near medium-sized stations or near small but highly trafficked stations, which are therefore particularly prone to congestion. In such situations, the blocked edges are typically surrounded by several potential alternative routes, which were not explicitly evaluated in this first simulation cycle. As a result, these cases should be re-examined through an additional round of simulations that explicitly incorporates re-routing options, where permitted by the network layout. In most cases, it is expected that the number of impacted services would decrease.

Nevertheless, some situations may arise – precisely those for which the proposed approach is most valuable – in which re-routing does not improve conditions and may even worsen overall performance. These cases indicate areas of the railway network that require further analysis also to assess the feasibility of re-routing trains along alternative paths not included in this first simulation cycle."

. Although a second simulation cycle was not performed within this deliverable, the results presented here were analysed on a case-by-case basis, considering both the location of the blocked edge and the feasibility of re-routing affected services.

While re-routing may introduce additional delays or congestion on alternative tracks and affect other services, it enables trains that would otherwise be completely blocked to complete their journey. Figure 32 highlights the cases with no feasible re-routing. These scenarios warrant further investigation, as they may address potential network improvements. While some of these cases are intuitive (e.g. station access points with only two tracks), others are less obvious and even counter-intuitive; addressing them may require targeted investments but will greatly improve the overall resilience of the infrastructure.

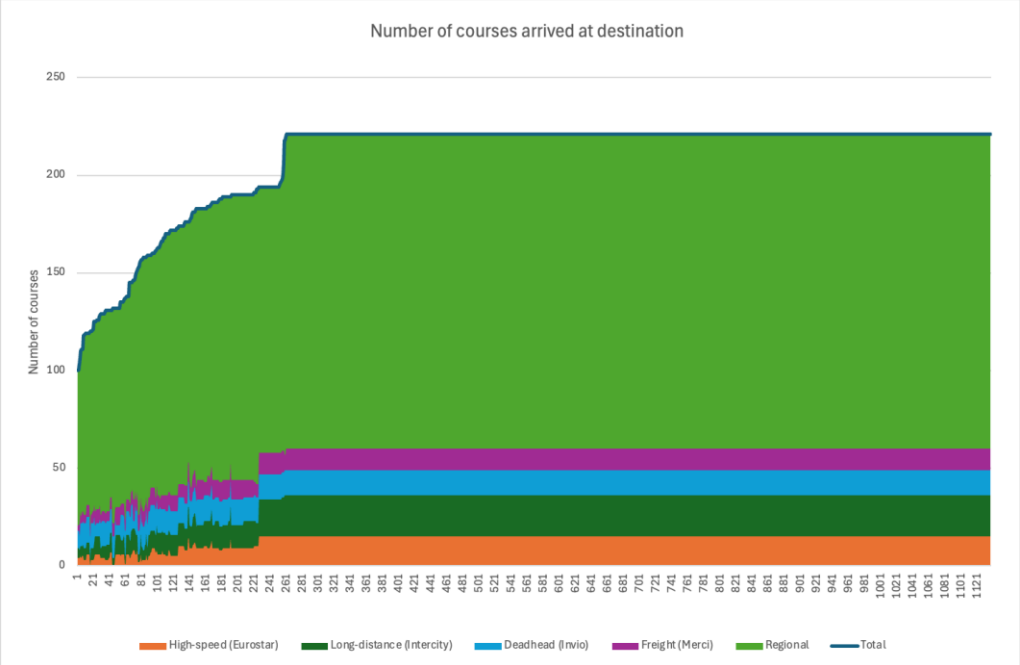

Figure 32: Courses arrived after re-routing evaluation

3.3 Findings and Discussion

In this section, the strengths and limitations of the adopted methodology are discussed, as they provide a useful basis for future developments.

First of all, the simulator selected for this study proved to be adequate – and ultimately the most appropriate choice – for the objectives of the project. The selection process required several meetings and in-depth discussions to evaluate and compare alternative tools, but in the end OpenTrack proved to offer the required modelling fidelity and operational realism for large-scale railway impact analysis.

At the same time, it should be acknowledged that OpenTrack is primarily designed as a planning tool and not to support such extensive, automated, and failure-oriented simulation campaigns. As a result, limitations emerged both in terms of performance and supported use cases, and – most importantly – in terms of missing programmatic controls.

Despite these constraints, the feasibility and robustness of the proposed methodology were demonstrated through the development of effective workarounds, enabling large numbers of simulations to be executed systematically. A residual simulator-related limitation concerns the simulation time window: courses whose scheduled departure falls outside the simulated interval are excluded from the analysis. This aspect becomes particularly relevant when adopting commonly used strategies such as running preliminary simulations only during peak hours to double check track occupancy and filter out less critical cases, followed by re-analysing the most significant scenarios over a full 24-hour horizon. In this context, restricting the simulation window may invalidate the results. Therefore, when dealing with fault scenarios, a 24-hours simulation horizon is strongly recommended.

Importantly, the software provider has also expressed availability to evaluate targeted improvements, informed by the experience gained within this project. Among the most important enhancements identified are:

- direct access to the routes traversing a given edge;
- the ability to set and modify train positions programmatically;
- more flexible programmatic management of routes;
- native support for signalling failures via the API; and
- improved and more expressive output formats, better suited for post-processing and resilience-oriented analyses.

Allowing the integration of more than one third-party tool within the SOAP channel would also be useful for further improving the simulation environment.

Despite these limitations, the methodology and the developed model offer several key advantages. The model is implemented programmatically, allowing extensive simulation campaigns to be executed by modifying only a limited set of parameters. Moreover, the use of the API, even with current constraints, enables failures to be injected automatically rather than through the graphical user interface, significantly reducing manual effort and enabling large-scale analyses that would otherwise be impractical. Through systematic simulations, the framework helps to identify the sections of the network that most significantly affects train operations and that would be difficult to detect using traditional approaches. In particular, the methodology enables the systematic analysis of all infrastructure components, without being limited to elements that are already known to be crucial, such as station approaches or single-track sections.

The approach developed is independent of both the size of the infrastructure and the specific railway section analysed, and is therefore fully transferable to other networks, ensuring its relevance well beyond the scope of this project. Naturally, as network size increases, simulation times also grow; exploring whether partitioning large networks into sub-networks – followed by the programmatic creation and management of such partitions – could accelerate simulations without compromising result quality represents a worthwhile direction for further investigation.

Although not strictly required, geo-referencing the model is strongly recommended. A geo-referenced model provides two major advantages:

- seamless integration of geo-referenced hazards, allowing failures derived from hazard scenarios to be incorporated into the OpenTrack model in a direct and precise manner;
- enhanced result assessment and aggregation, as simulation outputs can be associated with specific locations and aggregated at a geographical scale, facilitating the identification of macro-areas most affected by disruptions. This supports targeted interventions in areas with a high concentration of critical events, rather than acting on isolated and potentially distant points where mitigation efforts may be less effective.

The analyses presented correspond to a first simulation cycle, intentionally focused on severe scenarios – namely, complete blockages of infrastructure segments lasting an entire day and without considering mitigation measures. A second simulation cycle should revisit the most crucial cases by explicitly incorporating feasible re-routing strategies, which are expected to reduce impacts in many situations. Where no alternative routes exist, the results from the first cycle should be considered final. A third simulation cycle could further investigate failures involving railway switches, which may simultaneously block multiple directions and produce substantially more severe impacts.

Further developments should also address the sensitivity of results to disruption duration and time of occurrence, particularly for the most relevant scenarios. These aspects are essential to move from conservative worst-case assumptions toward more realistic, time-dependent impact assessments.

Additional developments could include the preventive processing of mitigation strategies—starting from re-routing policies and emergency plans—associated with each edge in the network graph and classified by expected duration and time of day. This would support the development of a ready-to-use decision-support system, helping to overcome the limitations imposed by long simulation times.

A challenging but highly relevant research direction concerns the optimisation of timetables, both under normal operating conditions and under disrupted scenarios. Investigating whether timetable optimisation – either preventive or adaptive – can be (pre)computed with the support of OpenTrack would represent a significant step toward advanced, simulation-driven decision support.

Finally, the results obtained should be interpreted as indications of potential areas of the network to be further assessed by domain experts. Expert evaluation is essential to assess the necessity of interventions, and define appropriate mitigation strategies. Such assessments should also consider the type of services affected and the corresponding priorities. While passenger services (high-speed, long-distance, regional) directly impact users, freight services may be equally or even more critical depending on the nature of the transported goods.

In conclusion, the proposed methodology and the tools developed in this work provide a robust and extensible framework to support railway operators. They enable the systematic assessment of impacts, the identification of potential structural vulnerabilities independently of specific hazards, and the planning of effective, proportionate mitigation strategies.

4. Roads

Traffic modelling plays a critical role in the design, offline evaluation, and real-time implementation of traffic surveillance and control strategies. Disruptions to the road network – such as bridge collapses or road damage – can significantly alter traffic conditions. In this context, traffic simulation tools are essential for assessing potential scenarios and supporting timely, informed decision-making.

This section focuses on the technical use of road traffic simulation to evaluate how disruptive events influence traffic distribution, rerouting patterns, and the restoration of network capacity. It supports resilience management by analysing the impacts of predefined network disruption scenarios. Traffic simulation models provide a robust analytical framework for evaluating network performance under such conditions. These models rely on origin–destination (O–D) flow matrices, which represent travel demand between zones, together with detailed network topology, link attributes, and performance functions. Traffic flows are assigned to the network based on travel times or generalised impedance across available alternative routes. This approach enables accurate estimation of prevailing traffic conditions and, when disruptions are introduced, supports the assessment of resulting changes in traffic distribution and accessibility.

The study is organized into three main phases: (i) mobility data analysis and (ii) development of a traffic simulation model for the study area, described in Section 4.1; and (iii) the design and evaluation of a set of prototype tests for possible events in the network, described in Section 4.2. Summary and discussion of key findings are presented in Section 4.3.

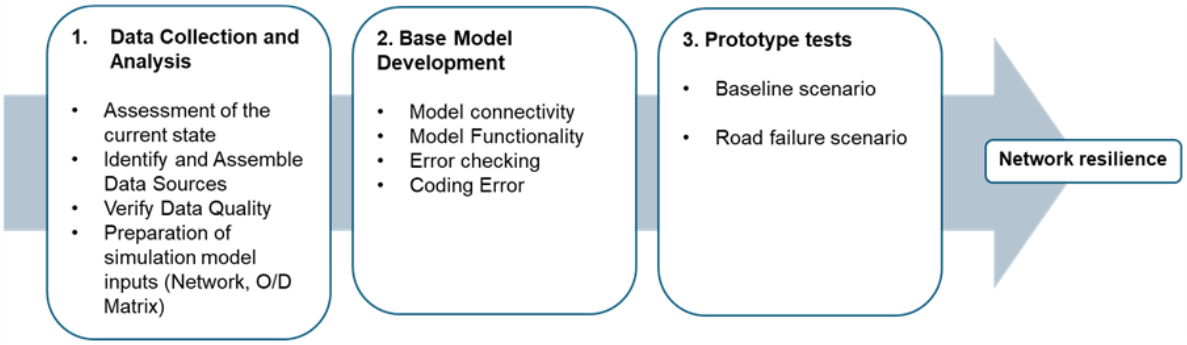

Figure 33: Workflow of the traffic modelling activity

The focus of this study is the strategic corridor involving the Pizzighettone bridge (see Figure 34), a critical infrastructure element characterized by high volumes of both heavy and light vehicles along the Milan–Marghera axis. Due to its role in supporting regional freight and passenger mobility, the bridge represents a key asset under both ordinary and disrupted network conditions and constitutes a central focus of complementary research activities within the RETURN project.

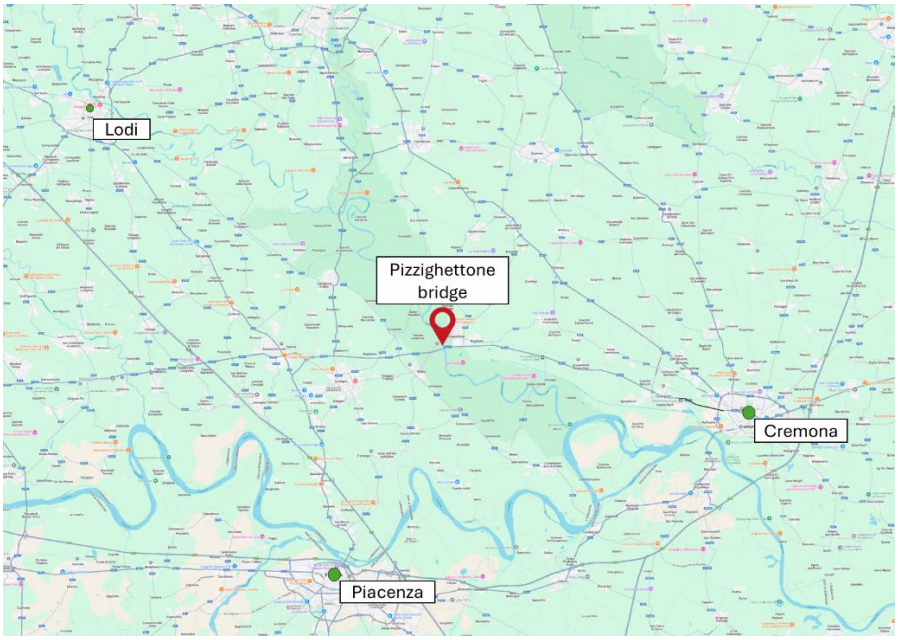

Figure 34: Study area

4.1 Methodology

4.1.1 Overview of simulation approach

Depending on the level of detail and computational granularity, traffic simulation models are typically classified into microscopic, mesoscopic, and macroscopic approaches (Hoogendoorn and Bovy, 2001; van Wageningen-Kessels et al., 2015).

Microscopic models simulate the behaviour of individual vehicles, explicitly representing their longitudinal dynamics (car-following) and lateral interactions (lane-changing) in response to surrounding traffic. This level of detail enables a realistic reproduction of driver behaviour and vehicle interactions, but it requires substantial computational resources, making microscopic models less suitable for large-scale networks or long simulation horizons.

At the opposite end of the spectrum, macroscopic models treat traffic as a continuous flow, analogous to a compressible fluid, and describe traffic dynamics using aggregated variables such as flow, mean speed, and density. These models are grounded in continuum flow theory and are particularly effective for large-scale strategic planning, where the focus lies on overall traffic patterns rather than individual vehicle behaviour. Nevertheless, their aggregated nature limits their ability to capture localized phenomena – such as shockwaves, bottlenecks, or incident impacts – in sufficient detail.

Mesoscopic models bridge the gap between microscopic and macroscopic approaches by combining aspects of behavioural realism with the computational efficiency of aggregate traffic flow formulations. In mesoscopic simulation, vehicle platoons are represented as discrete entities, while their movement along network links is governed by simplified – often analytical – performance relationships rather than detailed car-following rules. Traffic interactions, congestion formation, queue propagation, and spillback are captured through link performance functions and capacity constraints. This enables the simulation of dynamic traffic assignment and network loading in a computationally efficient manner, making mesoscopic models particularly well suited for evaluating traffic management strategies, incident scenarios, and resilience assessments in medium- to large-scale networks.

In this study, a mesoscopic modelling framework was selected, as it offers an optimal balance between behavioural detail and computational efficiency. The framework – originally developed by Sapienza University of Rome – has been adapted to the Milan–Marghera motorway corridor.

4.1.2 Data collection and pre-processing

Data collection and pre-processing typically represent the most resource-intensive stages of traffic analysis. Establishing a robust and reliable dataset – free from anomalies and inclusive of a representative range of traffic conditions – is essential to ensure the accuracy and validity of subsequent analytical and modelling activities. Large volumes of mobility data are collected daily from multiple tools and sources, and these datasets can serve several purposes, most notably as inputs for traffic simulation models and for assessing current traffic conditions across road networks of different scales.

In this study, data were sourced from the TomTom platform, which processes Floating Car Data (FCD) and GPS signals to generate traffic information. The platform provides multiple data formats, organized into categories such as traffic state indicators and origin–destination (O/D) matrices. A key advantage of this platform is its flexibility in defining the study area. Users can either draw a custom polygon or select predefined regional zones using the integrated GIS interface, with the corresponding road network being automatically identified and extracted. In addition, the platform allows the selection of a specific analysis period and supports temporal aggregation at 15-minute intervals.

However, converting raw data into application-ready inputs requires dedicated processing steps. The inputs of the traffic model developed in this study can be classified into two main groups: *static* and *dynamic*. In both cases, the raw data obtained from the TomTom platform required manual processing and adaptation before being integrated into the model.

Static inputs describe the network's topological, geometric, and functional characteristics and include all elements required for realistic traffic simulation. These comprise lane configurations, free-flow speeds, emergency lanes, toll stations, bridges, viaducts, speed limits, and model parameters that remain constant over time. In accordance with the Smart Road Decree, each road is modelled as one or more arcs with nodes at their endpoints. Furthermore, motorway sections (excluding ramps) are subdivided by fictitious nodes whenever changes in functional characteristics occur.

Dynamic inputs, by contrast, vary over time and include elements such as the O–D matrix and traffic measurements. These data are essential not only for executing traffic simulations but also for model calibration and for improving overall accuracy.

As mentioned, the mobility database used in this study is based on Floating Car Data (FCD) provided by TomTom. This information is collected from multiple sources, including vehicle GPS devices, mobile operating systems (iOS and Android), fleet management systems, and proprietary sensors. The data are continuously archived and processed, then map-matched to the TomTom digital map either in real time or during report generation. Two types of data are particularly relevant for this study:

- Origin-Destination (OD) data, providing the volume of movements between locations (at a spatial resolution not finer than the municipal level), and
- Map matched traffic state data, , including average speed and travel time of sampled vehicles, together with the corresponding sample size for each segment of the selected network.

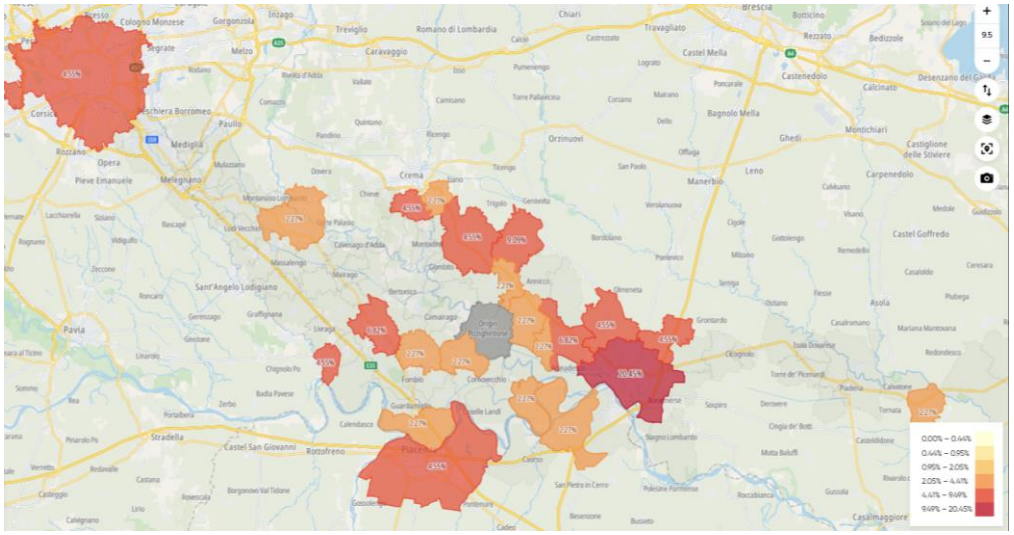

(a)

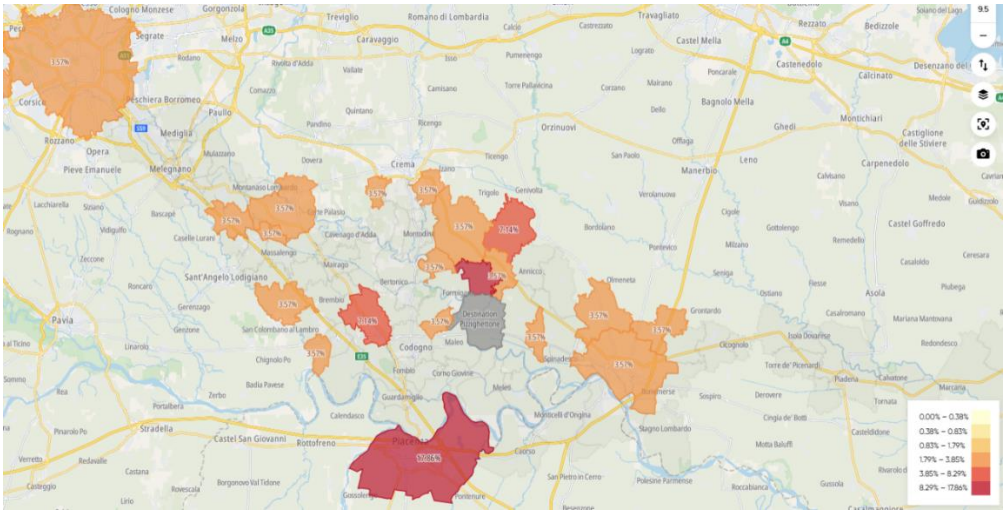

(b)

Figure 35: Visualization of OD function of TomTom platform, a) Percentage distribution of trips by destination municipality with Pizzighettone as the origin, and b) Percentage distribution of trips by destination municipality with Pizzighettone as the destination

Figure 36: Visualization of traffic state function of TomTom platform

Network graph

The network analysed in this study covers a wide area, including primary and secondary expressways as well as several urban road segments between Lodi and Cremona (both located in central Italy). This corridor is strategically important, as it includes the Pizzighettone bridge crossing the Adda River.

Before applying the model, the road network was converted into a link–node representation incorporating the essential road characteristics required as model inputs. In this structure, each continuous road segment is modelled as a link up to the point where it intersects with another road or connects to an entry/exit ramp. The beginning and end of each link are represented by nodes.

Following this discretisation process, the resulting network – shown in Figure 37 – consists of 9.650 links and 6.285 nodes.

Figure 37: Link-Node graph of study area

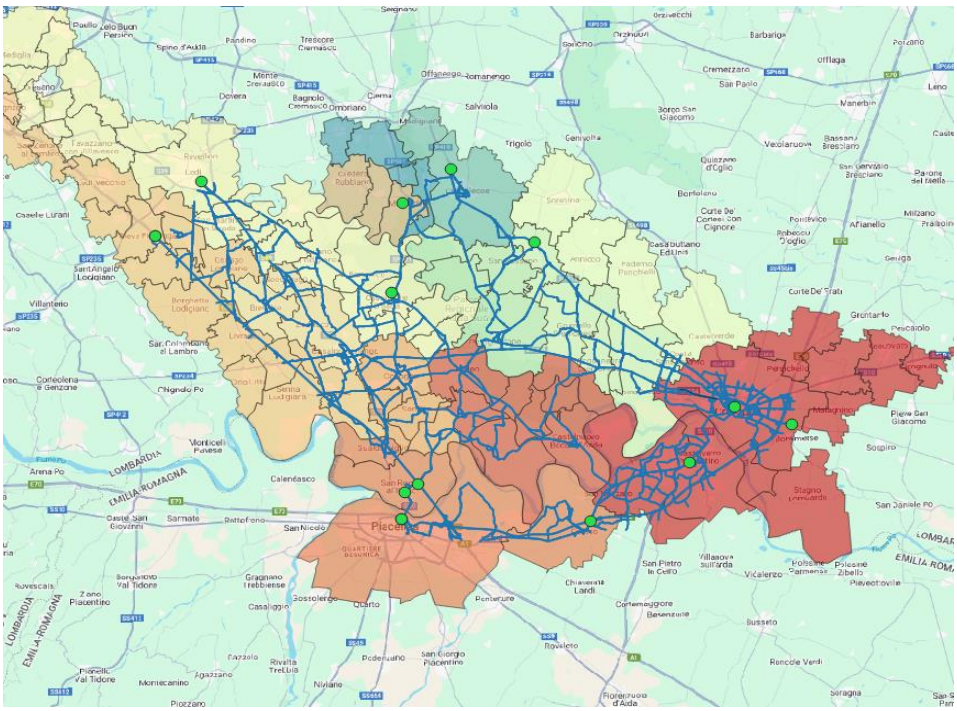

Figure 38: Network, zoning and centroids visualizations

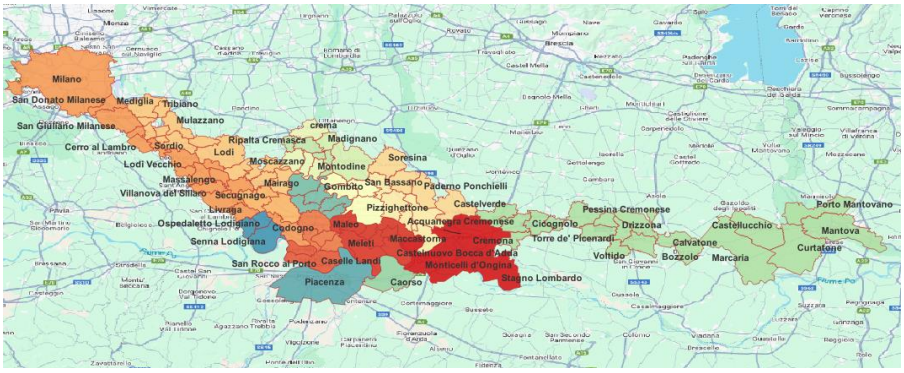

Figure 39: Zoning of study area

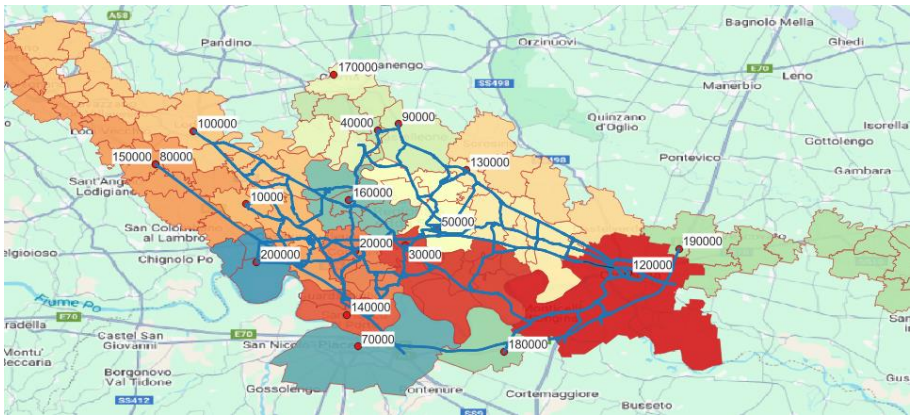

Figure 40: Centroids of zones

The preparation of the network for the simulation model involved several key steps. First, the study area was divided into traffic analysis zones by aggregating nearby areas into single zones, each represented by a centroid to simplify demand representation. In total, 20 zones were defined, as illustrated in Figure 39. The network links and nodes were configured to ensure full connectivity between all centroid pairs.

Essential attributes required by the model – such as the number of lanes, maximum speed (Figure 41), and the parameters of the fundamental diagram, which describe the relationships among traffic variables for each network section – were manually assigned. Additional adjustments, including the specification of ramp links and the modification of connectors, were also performed at this stage to ensure an accurate representation of the network.

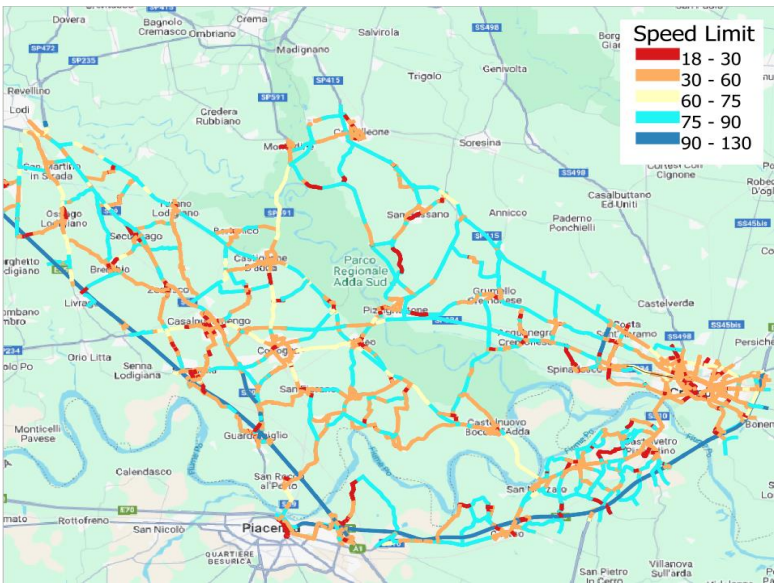

Figure 41: Extended network's speed limits

Table 3 lists the link attributes included in the shapefile and their role in the traffic simulation model.

Table 3 – Graph's attributes

Graph attribute	Note
fl_id	Link ID
fl_from	Initial node of link
fl_to	Final node of link
fl_v0	Free flow speed (km/h)
fl_length	Length (km)
fl_capacity	Capacity (veh/h)
fl_alpha	Exponential coefficient of fundamental diagram
fl_r_cr	Critical density (veh/km)
fl_lanes	Number of lanes
fl_connector	Flag connector links

Traffic state

Figure 36 illustrates the traffic conditions along the route from Milan to Cremona, passing through Pizzighettone. The average vehicle speed, derived from collected Floating Car Data, is represented using a colour scale ranging from yellow to red, where warmer colours indicate higher levels of congestion or reduced speeds.

In accordance with the road characteristics, the carriageway configuration (divided or undivided) and the applicable speed limits were used to classify the road network and assign appropriate fundamental diagrams (FD). These fundamental diagrams reflect the specific operational characteristics of each road category and form the basis for modelling traffic flow dynamics within the simulation.

Figures 42, 43, and 44, together with Table 4, illustrate respectively the spatial distribution of the adopted fundamental diagrams, the flow–density relationship, the speed–density relationship associated with each FD configuration, and the corresponding parameters adopted for the study network.

Figure 42: Fundamental diagram map

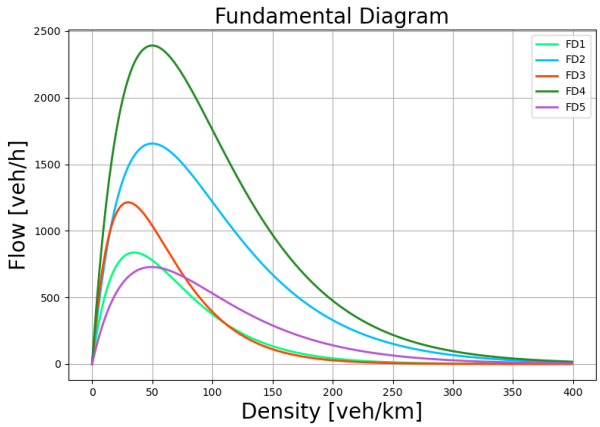

Figure 43: Flow-density relationship for each set of fundamental diagrams

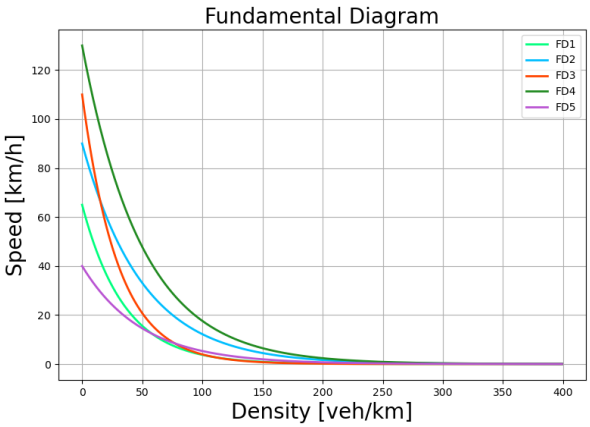

Figure 44: Speed density relationship for each fundamental diagram

Table 4 – Fundamental diagram values

FD N.	Street code	Vmax	rcr	a
1	Municipal road	60	30	1
2	Provincial Street	90	50	1
3	Separate carriageway	110	30	1
4	Highway	130	50	1
5	Ramp	40	30	1

Traffic conditions along the Milano–Cremona corridor were analysed in terms of the number of samples and the corresponding average speed over the entire route. The analysis was performed for each day of the week over a one-month period, focusing on the morning peak hours (06:00–08:00). As shown in Figure 45, Monday recorded the highest number of samples, in line with typical commuter traffic patterns, while Saturday and Sunday exhibited markedly lower sample volumes due to reduced traffic demand. The remaining weekdays displayed comparable sample levels, indicating relatively stable travel demand throughout the working week. The analysis of average speeds further confirms this pattern: weekday speeds were relatively consistent and slightly lower as a result of higher congestion levels, whereas weekend speeds were significantly higher, reflecting lighter traffic conditions.

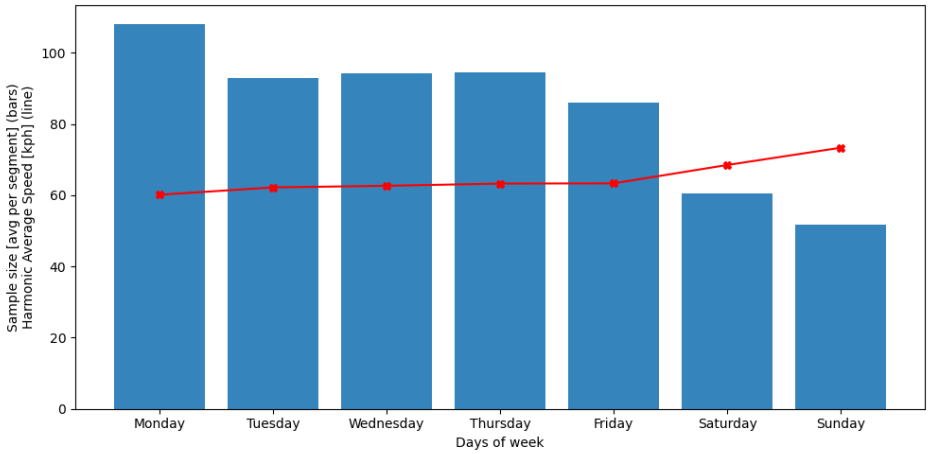

Figure 45: Assessment of traffic conditions along a specific route during the first week of August 2024 using TomTom data.

O-D Matrix

The Origin-Destination (O-D) matrix, representing movements between municipalities at 15-minute intervals, is derived from TomTom data in the form of a matrix in which rows correspond to origins and columns to destinations. This matrix was subsequently reformatted to match the input requirements of the traffic simulation model. Figure 46 provides an immediate overview of the magnitude of displacements between the different locations.

Figure 46: Desire lines representing the movements between two places, darker color shows higher number of displacements

The analysis of the number of trips passing through the study area in 2025, aggregated over four months and shown in Figure 47, reveals clear seasonal and temporal variations in traffic volumes. The monthly aggregation indicates that October records the highest mobility level, exceeding 25 million trips, followed by May, with approximately 20 million trips. Trip volumes decrease noticeably during the summer period, reaching a minimum in July, which highlights a general seasonal downturn. This pattern reflects the combined effects of seasonal travel behaviour, commuting intensity, and holiday periods along an extended corridor where both expressway and urban road segments contribute to the overall traffic volume.

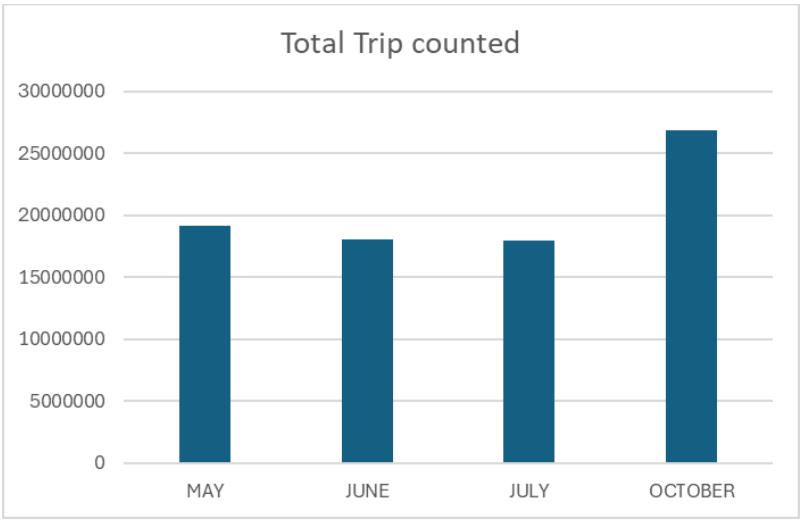

Figure 47: Total number of trips counted for the wide network in 4 months (May, June, July, and October)

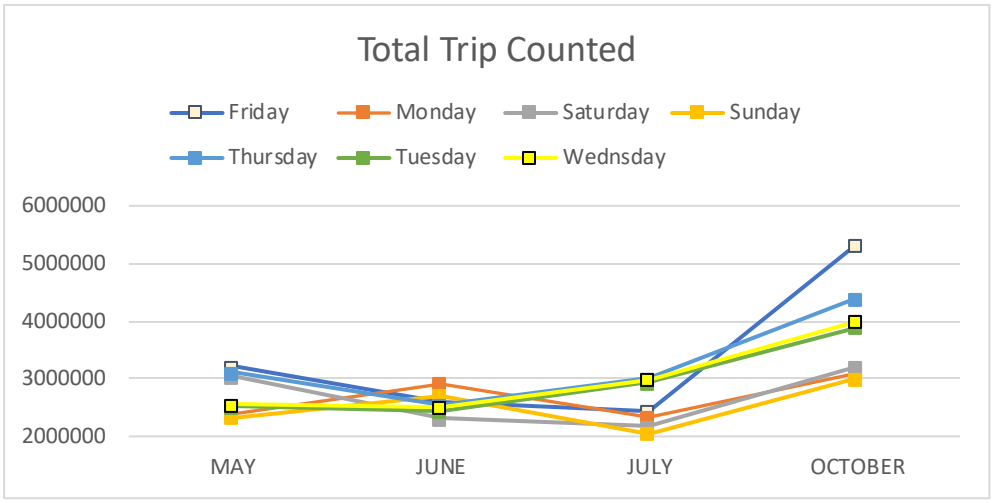

Figure 48: Total number of trips counted on each day of week for 4 months (May, June, July, and October)

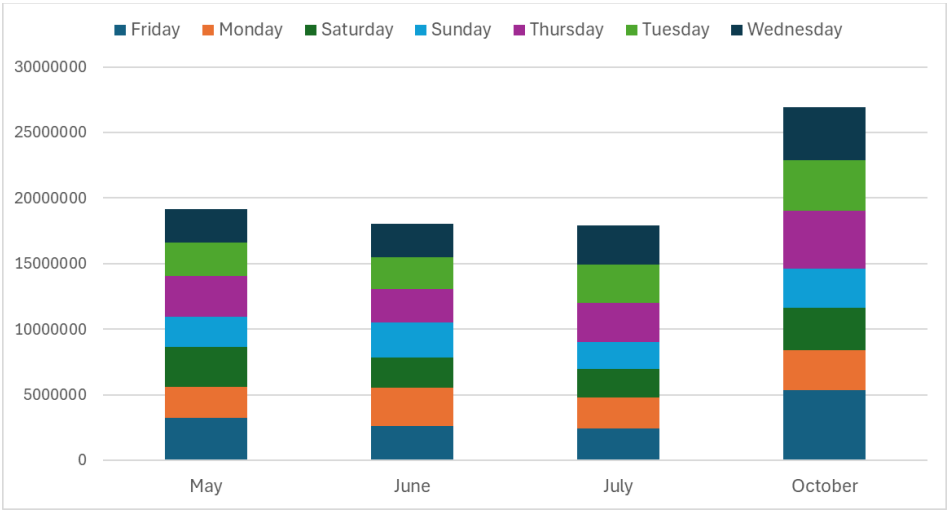

Figure 49: Total number of trips in the study area by day of the week over four selected months (May, June, July, and October), showing the distribution of travel demand across weekdays and weekends.

Examining the distribution of trips by day of the week for each month (Figures 48 and 49) reveals more detailed temporal patterns. Weekdays generally exhibit higher demand than weekends, consistent with work-related mobility across an extensive road network. Among all combinations, Friday in October stands out with the highest number of trips, exceeding 5.4 million, indicating a pronounced peak in weekday traffic intensity during that month. Conversely, October also shows the largest weekday–weekend contrast, with Sunday experiencing a marked reduction in trips.

As the summer period progresses into July and later into October, variations among weekdays begin to smooth out, reflecting reduced commuting flows and more uniform travel behaviour throughout the week. Together with Figure 47, these graphs illustrate how both monthly and weekly dynamics shape traffic demand across a large regional road network.

Figure 50 illustrates the temporal evolution of total daily trips for each day of the week from May to October, providing a detailed view of how trip volumes fluctuate over time and across months. The line chart shows that Thursday and Friday consistently exhibit the highest trip volumes, while Saturday and Sunday remain

significantly lower throughout the period. Notably, Thursday maintains trip levels comparable to Friday but without the higher variability and early-weekend effects that often characterise Friday travel patterns.

Considering the overall trends observed across the analysed months, the first or second Thursday of October emerges as the most representative day for traffic analysis. During this period, total trip volumes reach their maximum, indicating a return to full activity levels following the summer slowdown. Thursday consistently reflects typical weekday travel behaviour, combining high demand with stable and regular patterns. Selecting an early October Thursday therefore ensures the capture of representative weekday conditions under full traffic activity, making it the most suitable choice for the analysis.

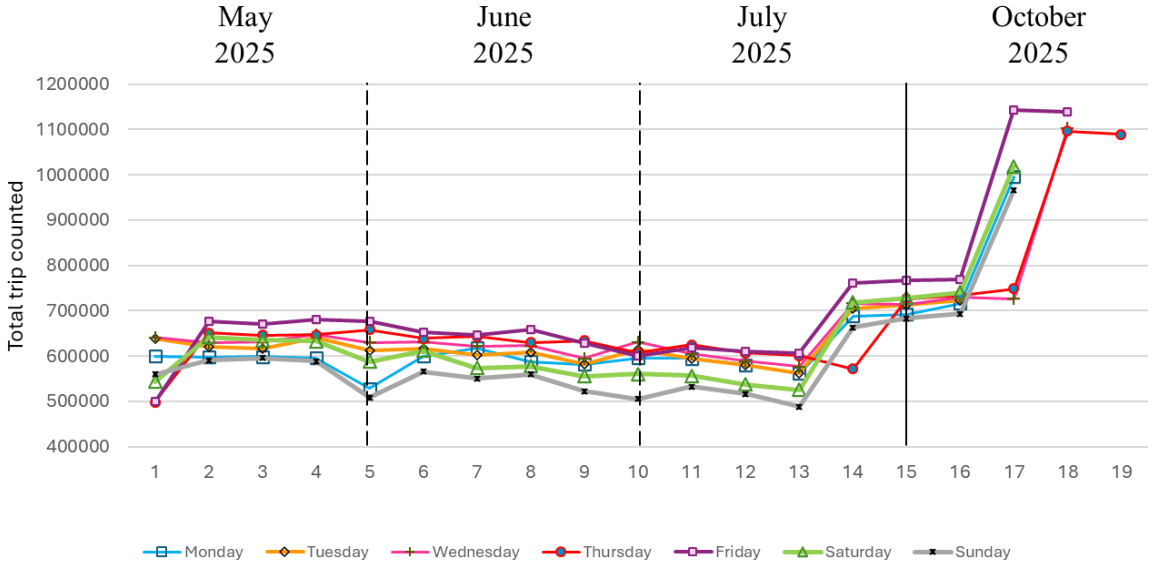

Figure 50: Variation of total daily trips by day of the week from May to October, illustrating temporal changes in travel activity and weekday trends over the study period.

4.1.3 Model Development Steps

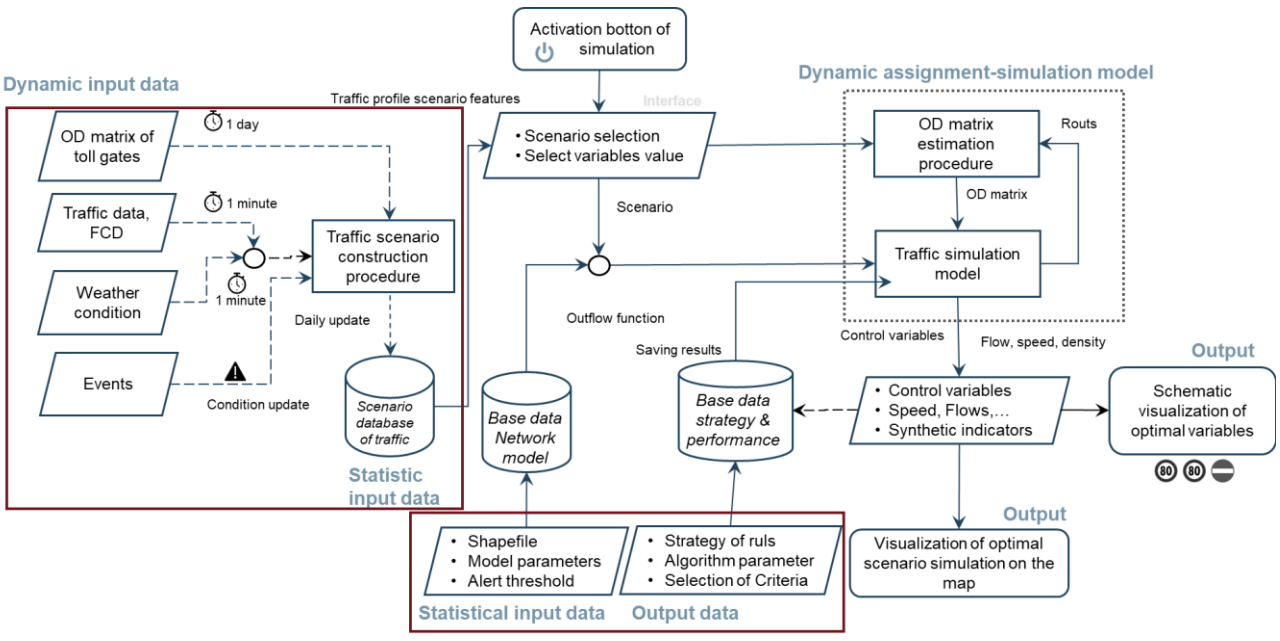

Figure 51: Traffic simulation framework

The traffic simulation model adopted in this study is a mesoscopic dynamic flow framework developed by Sapienza University of Rome within the MOST project. The model was specifically designed as a decision-support tool for traffic management, offering a high degree of flexibility in modifying both supply-side and demand-side parameters. It allows the application of network-wide or localised changes to road infrastructure and operating conditions and dynamically simulates the resulting traffic flows across the network.

The simulator supports a wide range of scenario-based analyses, including partial or complete closure or reclassification of road segments; the application of adverse weather conditions (rain, wind, snow) with three predefined levels of severity; and the restriction or exclusion of specific vehicle classes – such as heavy goods vehicles – either on selected links or across the entire network. These features make the model particularly suitable for evaluating resilience strategies, incident response measures, and emergency traffic management policies.

To improve the computational efficiency of the dynamic equilibrium assignment procedure, the model avoids repeating the computationally intensive shortest-path search at each iteration. Instead, it precomputes the K optimal paths between each origin–destination pair (forming a forest of K optimal paths) and updates only the travel times associated with the previously identified paths.

Traffic dynamics are simulated by assigning each vehicle to a route that is dynamically updated based on prevailing network conditions at the time of departure. Vehicle speeds are computed as a function of the average density within each traffic loop, vehicle positions are updated at 6-second intervals, and movements at loop exits follow a First-In–First-Out (FIFO) principle. When traffic demand exceeds either the flow capacity of a loop or the storage capacity of a downstream segment, queues are formed and explicitly tracked.

Node regulation is represented through traffic signal control, modelled using queue formation and discharge processes, as well as through precedence rules implemented via gap-acceptance behaviour. The effects of heterogeneous vehicle performance and driver behaviour are captured through differentiated coefficients defining the speed–density performance functions on traffic loops and behavioural rules at network nodes.

4.2 Simulations and Results

4.2.1 Prototype scenario design

The testing and validation of the dynamic network simulation–assignment procedure were carried out by applying the developed model to the entire study area and comparing the simulation outputs with the available traffic data at the required level of detail, with particular reference to the Pizzighettone bridge.

The validation procedure was structured into two main categories:

1. Network connectivity tests, consisting of elementary simulations designed to verify network connectivity, route-choice behaviour, and the model's ability to reproduce traffic dynamics under simplified abnormal conditions, using a single origin–destination (O–D) pair.
2. Model functionality tests, conducted using the full O–D matrix estimated from Floating Car Data (FCD), in order to evaluate model performance under realistic demand levels and to assess its capability to dynamically reassign flows and explore alternative routes across the extended network.

The two categories of tests described above were further articulated into the following specific test cases:

- Network connectivity tests

  a. Dynamic assignment test with a single O–D pair
  b. Lane closure test (one lanes)
  c. Carriageway closure test

- Model functionality tests

  a. Dynamic assignment test with the estimated O–D seed matrix
  b. Carriageway closure test with the estimated O–D matrix (bridge closure)

The accuracy test evaluates the model's ability to reproduce real traffic conditions using the available input data by calculating a set of error indicators based on speed measurements.

- Mean Absolute Error (MAE):

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|\hat{x}_i - x_i^o|$$

- Root Mean Square Error (RMSE):

$$RMSE = \left[\frac{1}{n}\sum_{i=1}^{n}(\hat{x}_i - x_i^o)^2\right]^{\frac{1}{2}}$$

where:

- x_i^o is the generic observation of the variable x
- $\hat{x}_i$ is the model estimate of the variable x corresponding to the i-th observation
- $\underline{x}$ is the meaning of the n observed values

4.2.2 Network connectivity tests

The basic functionality tests – identified as Test a. Base, Test b. Lane closure, and Test c. Road failure – were performed in simulation mode using a single demand loading on the shortest path (MAX_ITE = 1). These tests were designed to verify the expected changes in simulation outputs resulting from variations introduced in the model input data, specifically modifications to road network attributes. The results were then compared against a reference scenario in which no attribute modifications were applied.

Test a. Base

The demand matrix consists of a single origin–destination (O–D) pair, Lodi–Cremona, with a flow of 1.000 vehicles per hour over the time interval 07:00–09:00. Table 5 and Figure 52 show the location of the corresponding centroids. Figure 20 illustrates the simulated movement of the vehicle platoon through the network, from centroid 100000 (Lodi) to centroid 120000 (Cremona), starting at the beginning of the simulation.

This simplified demand configuration ensures that the effects observed in the simulations can be directly attributed to changes in network attributes, enabling a clear and controlled evaluation of the model's response.

Table 5 – O-D matrix of single platoon

O	D	Classe	Val	ts	Timestamp
100000	120000	0	1000	435	07:15
100000	120000	0	1000	450	07:30
100000	120000	0	1000	465	07:45
100000	120000	0	1000	480	08:00
100000	120000	0	1000	495	08:15

Figure 52: One pair of O-D for model functionality test

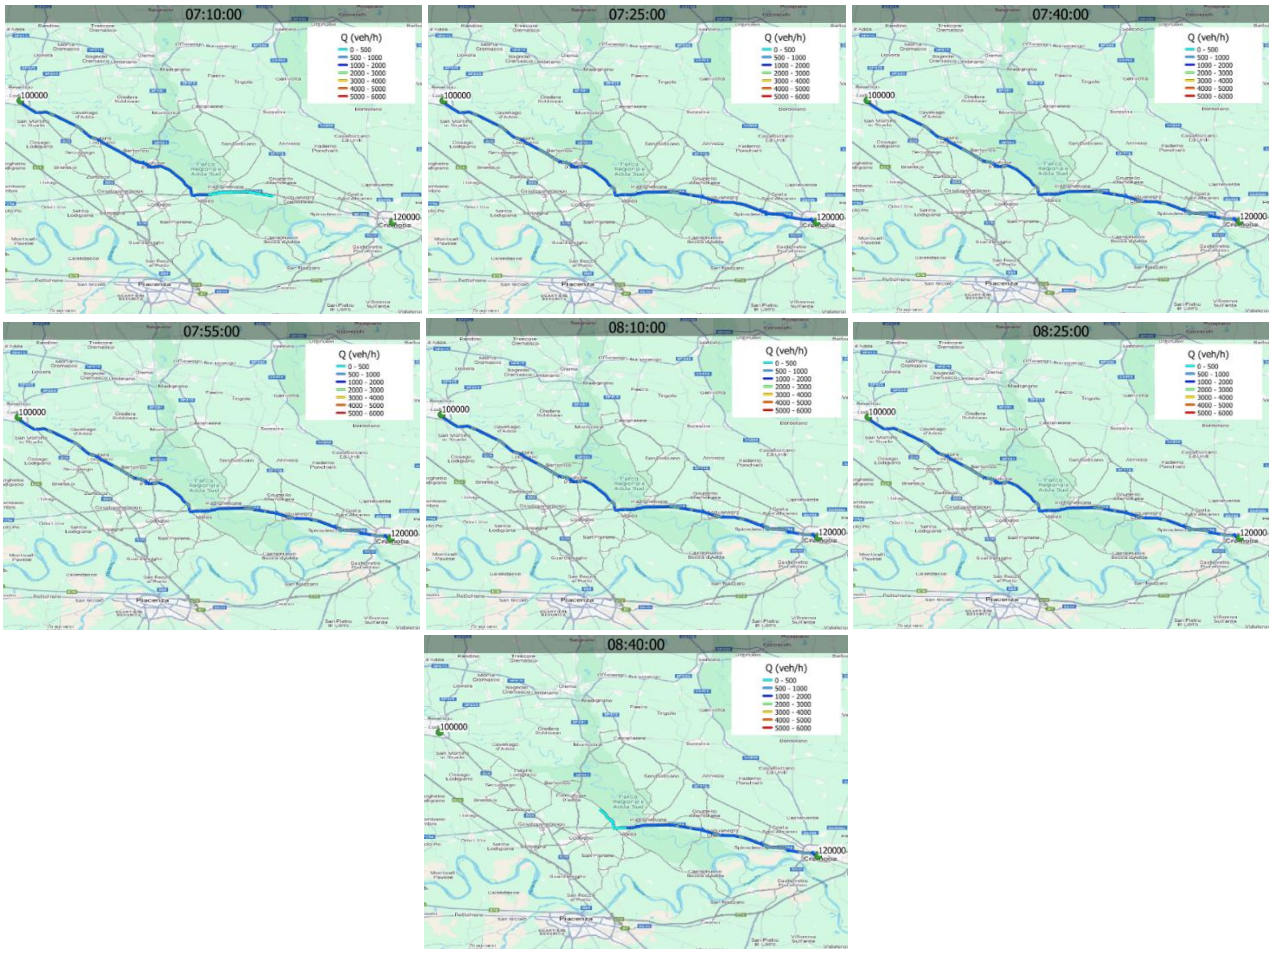

Figure 53: Test a. results

The test expectations for the single origin–destination pair are met, as the model successfully demonstrates network connectivity from the westernmost to the easternmost part of the study area across the extended network (Figure 53). The simulation accurately captures the progressive build-up of traffic flow over time, followed by a gradual discharge of vehicles from the origin toward the end of the simulation period.

Test b. Lane closure

In this test, a single demand platoon is feeded into the network, and the Pizzighettone bridge with the highest traffic demand is closed, while an alternative bridge with a lower speed limit remains open to accommodate rerouted traffic.

Figure 54: Pizzighettone bridges

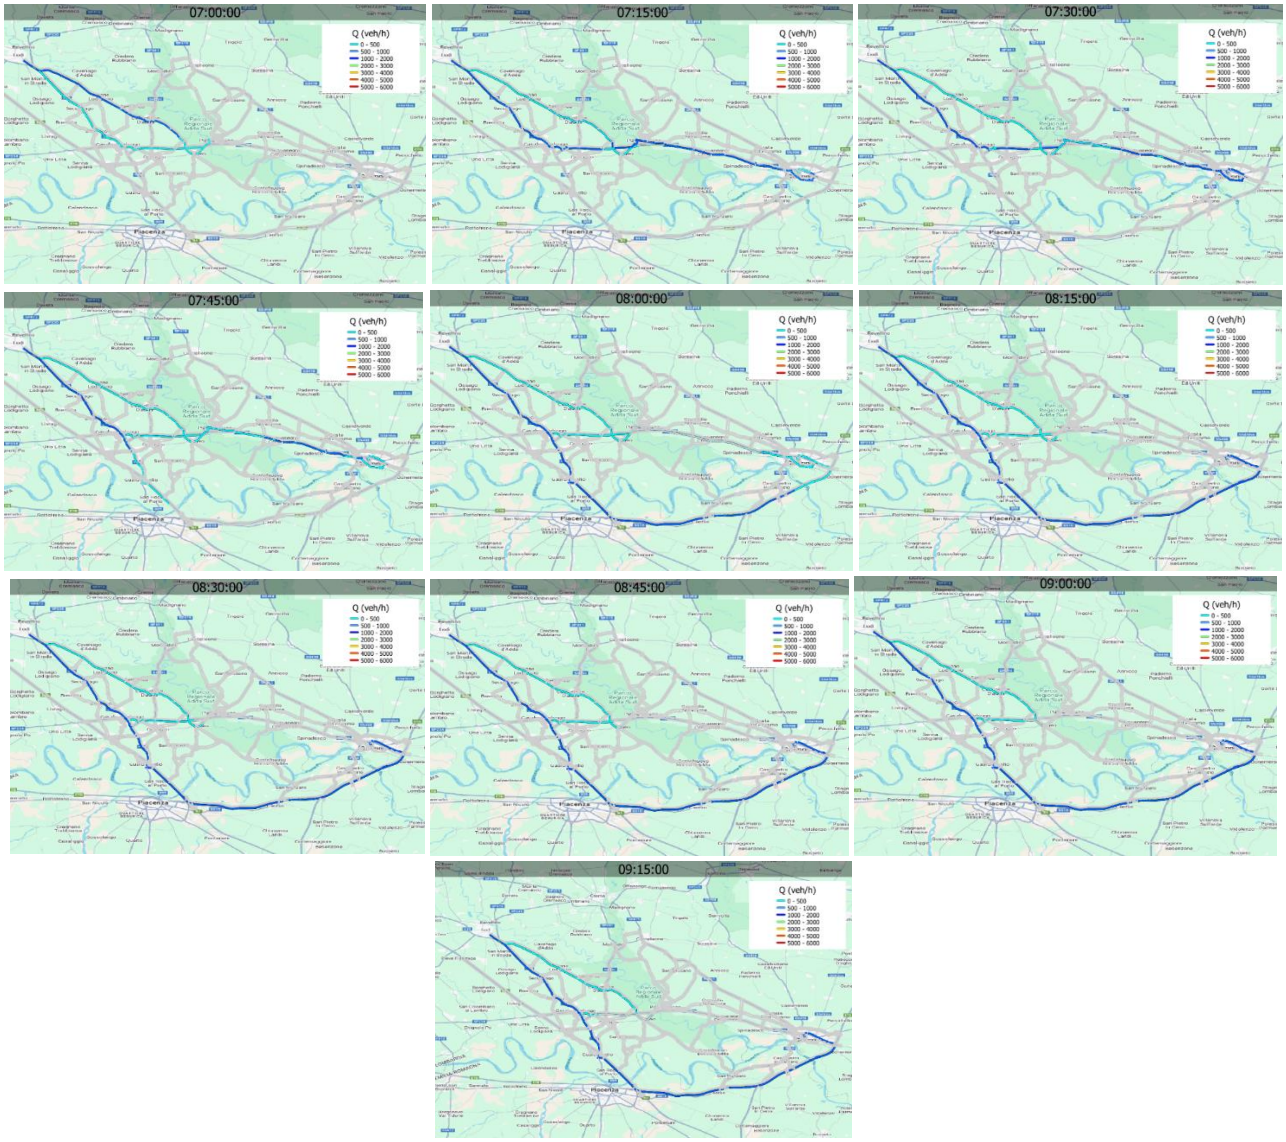

Figure 55: Test b. results

The objective of this test is to verify the correct connectivity of the network graph and the proper functioning of the model in reproducing traffic flow responses to a road closure along the path between a single origin–destination pair. The expected outcome is the formation and upstream propagation of congestion from the closure point, accompanied by the rerouting of traffic toward available alternative routes.

As illustrated in Figure 55, once the closure is applied at 07:20, the simulation shows a clear accumulation of demand upstream of the blocked section. Initially, the affected flow is rerouted upstream, making use of the adjacent bridge near Pizzighettone around 07:30. Subsequently, as congestion builds up and alternative routes become more attractive, a progressive redistribution of traffic is observed, with an increasing share of demand shifting toward the faster route via the highway by approximately 07:45.

Test c. Road failure

In this test, a single demand platoon is introduced into the network, and both bridges in Pizzighettone (Figure 54) are closed from 08:00 to 08:45, while the simulation covers the time interval 07:00–09:00. The objective of the simulation is to evaluate the model's ability to dynamically adapt to disrupted network conditions and to identify alternative routes that still allow the demand to reach its destination.

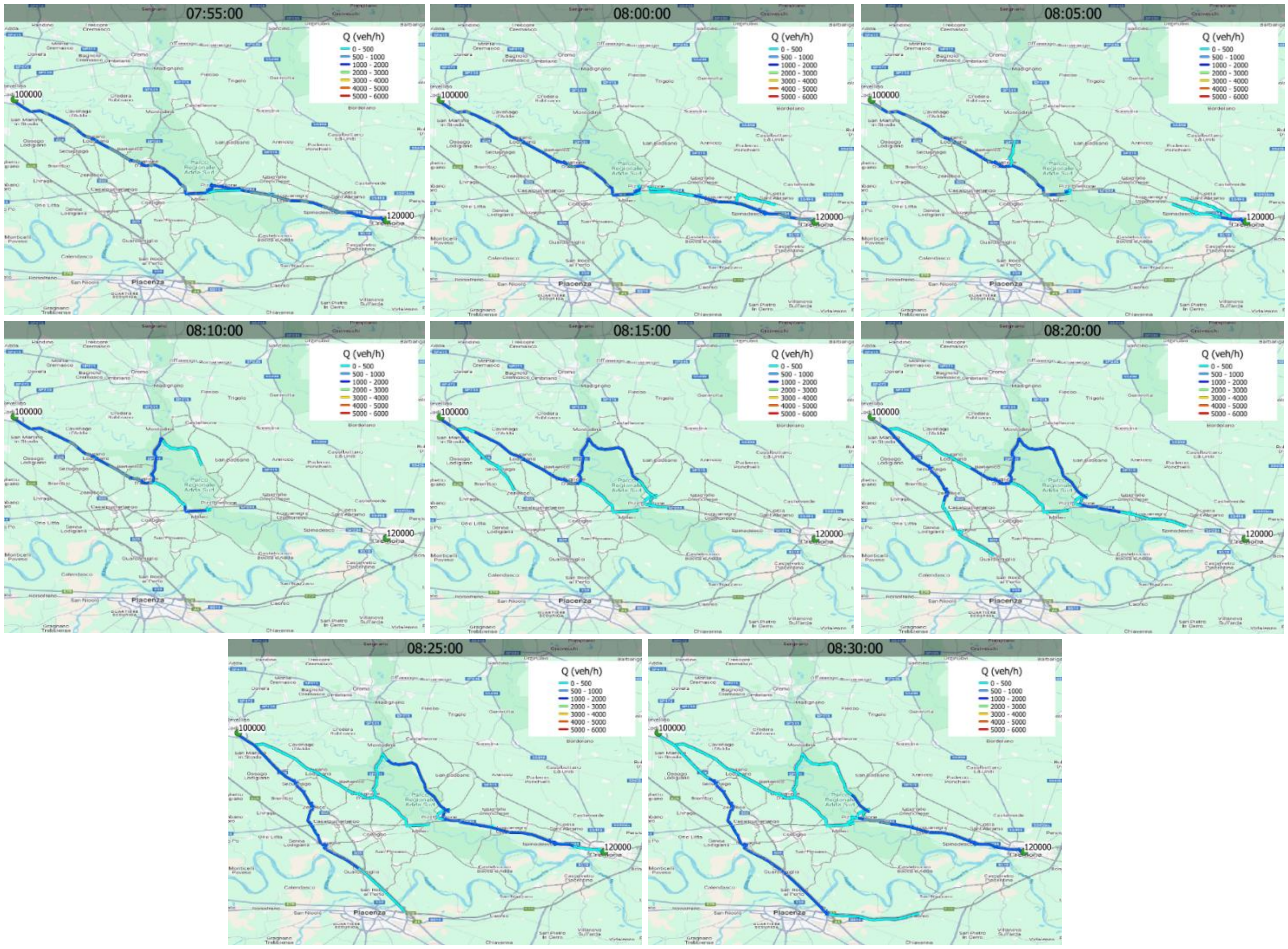

Figure 56: Test c. results

Figure 56 illustrates the simulated traffic conditions under normal operating conditions, where the demand platoon follows the shortest path and traffic flow remains stable across the network. When the bridge closure is introduced during the predefined time interval, the model dynamically detects the disruption and reallocates traffic flow onto the available alternative route. This rerouting mechanism preserves network continuity and enables the demand to reach its destination despite the infrastructure failure, demonstrating the model's ability to reproduce adaptive traffic behaviour under critical conditions.

4.2.3 Model functionality tests (Designed scenarios)

The Pizzighettone Bridge is located along the SP234 road, which serves as the designated route for authorised heavy vehicles. The following analyses assess the extent of changes in traffic distribution that would occur in the event of a bridge collapse or a temporary loss of functionality.

For these tests, the origin–destination (O–D) seed matrix corresponds to the complete dataset obtained from the TomTom platform, aggregated according to the zoning system defined and described in the previous sections. To extract the O–D matrices, an external region analysis approach was adopted. This method quantifies trips that originate within, terminate within, or pass through the predefined study area. In addition, it separately accounts for trips that originate outside the study area but either end within it or traverse it, as well as trips that originate within the study area and either end outside it or pass through the external region.

Figure 57: External region trip counting concept

The analysis records the number of trips between each pair of municipalities within the study area. It also accounts for trips that:

(i) originate outside the selected municipalities and terminate within the study area;

(ii) originate within the external region and terminate outside the study area; and

(iii) both originate and terminate outside the study area but pass through the defined area.

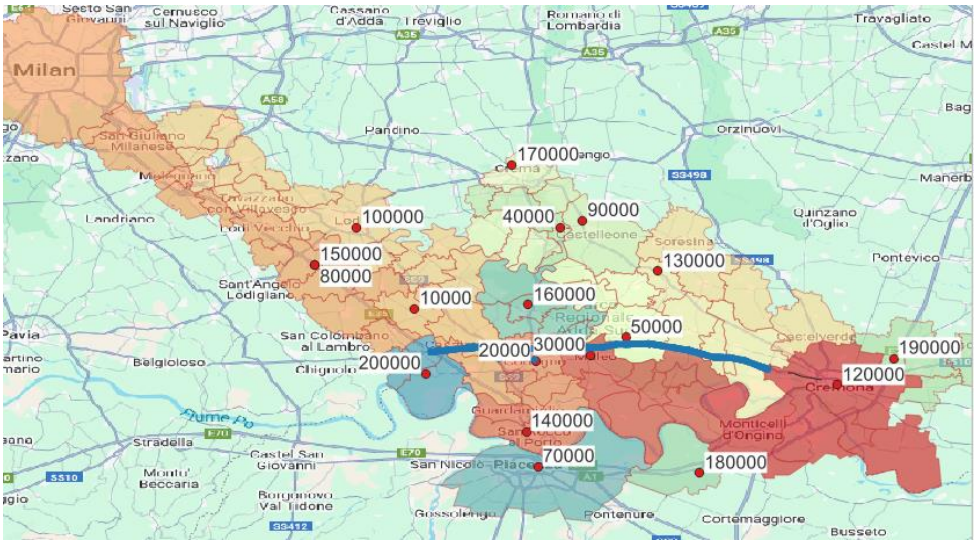

Figure 58: SP234

To capture the impact of trips that use the network without clearly identifiable origins or destinations (i.e. pure through traffic), an additional modelling assumption was introduced. Specifically, for trips that originate within the study area but whose exact destinations cannot be reliably determined, destinations were assumed to correspond to one of three major regional attractors surrounding the study area. These attractors

were defined as Milan, Cremona, and Piacenza, representing the northern, eastern, and southern directions, respectively.

This assumption enables the inclusion of network load effects generated by through and outbound traffic, ensuring that the contribution of such movements is adequately represented in the O–D matrix estimation and in the subsequent network analysis.

	Destination																	
Origins	10000	20000	30000	40000	50000	70000	80000	90000	100000	120000	130000	140000	160000	170000	180000	190000	200000	generation
10000		26	2	0	1	27	51	2	15	23	1	3	3	1	0	0	2	157
20000	29		3	0	2	37	43	1	12	25	0	5	5	0	1	0	7	171
30000	3	4		0	4	11	5	0	5	9	0	1	0	0	0	0	0	43
40000	0	0	0		0	15	16	2	7	15	0	0	0	56	0	1	0	111
50000	2	9	5	2		9	10	0	2	22	0	2	0	3	0	3	0	70
70000	36	40	7	12	7	841	6712	22	385	633	15	34	7	352	41	23	20	9186
90000	0	0	0	0	1	20	26		4	38	5	0	0	88	0	1	0	184
100000	11	9	0	0	0	222	742	4		219	0	5	1	27	0	1	1	1241
120000	32	41	4	13	15	787	6587	27	378	517	21	25	8	355	30	45	17	8901
130000	0	1	0	0	2	15	16	9	0	48		0	0	25	0	1	0	117
140000	8	10	2	0	0	74	23	1	2	15	1		0	1	2	0	0	138
150000	60	51	4	13	7	3511	9289	24	691	3261	18	32	9	366	19	23	33	17410
160000	10	14	0	0	1	9	10	1	11	7	0	2		4	1	0	2	72
170000	0	0	0	12	1	99	148	20	18	107	2	0	2		0	0	0	410
180000	0	0	0	0	0	52	18	0	0	19	0	1	0	0		0	1	92
190000	0	0	0	0	1	25	27	1	1	123	0	0	0	1	1		0	179
200000	22	17	0	0	0	14	26	0	7	12	0	1	1	0	1	0		101
Attraction	213	222	27	52	41	5768	23749	113	1538	5093	62	110	36	1280	97	99	82	38583

Figure 59: Origin-Destination matrix for hour 7 to 8

Baseline Scenario

The Baseline Scenario represents the normal operating conditions of the transport network, without any disruptions, restrictions, or infrastructure failures. It serves as the reference configuration for evaluating system performance under ordinary demand and supply conditions, based on the available traffic data, calibrated model parameters, and validated demand matrices.

A comparison between the observed Floating Car Data (FCD) (Figure 60) and the simulation results (Figure 61) shows a moderate discrepancy in the estimated speed and travel time for the first bridge segment. For the bridge entrance segment (length = 70 m), the observed data indicate an average speed of approximately 63 km/h (corresponding to a travel time of 73 s), whereas the simulation reports an average speed of about 89 km/h (approximately 55 s travel time). This difference suggests that the model represents slightly smoother traffic conditions on this short link.

Given the limited length of the segment and the inherent variability associated with floating-car sampling, such deviations are not unexpected and do not indicate a significant mismatch. Nevertheless, the comparison highlights an opportunity for minor calibration refinements to further improve the alignment between observed and simulated performance, particularly in areas characterised by short links and localized speed variations.

The comparison between observed and simulated speeds indicates that the model is able to reproduce the overall level and variability of traffic speeds, albeit with non-negligible errors. An MAE of approximately 18 km/h and an RMSE of about 22 km/h suggest that, on average, simulated speeds deviate from field measurements by roughly one speed class, with larger discrepancies occurring on some segments and during specific time periods.

While this level of accuracy is not sufficient for detailed operational assessments—such as precise estimation of queue lengths or level-of-service at individual links—it is acceptable for exploratory analyses and scenario testing, particularly when the objective is to evaluate relative changes between scenarios rather than exact absolute values.

The validation of the traffic simulation results focuses on the road segments located in proximity to the Pizzighettone bridge, where reliable external traffic information is available. For this purpose, TomTom

traffic state data were used, providing the average current speed on each road segment and enabling comparison with the corresponding free-flow speed. As shown in Figure 63a, the traffic state observed at the selected locations along the road segments approaching the bridge indicates that current speeds closely match free-flow conditions.

Figure 63b presents the simulated speed results using a colour-coded representation. Simulated speeds range between 80 and 100 km/h, in accordance with the applicable road regulations and closely aligned with the speeds reported by TomTom. This level of agreement confirms that the model is able to realistically reproduce prevailing traffic conditions in the analysed area, thereby supporting the validity of the simulation results for the considered baseline scenario.

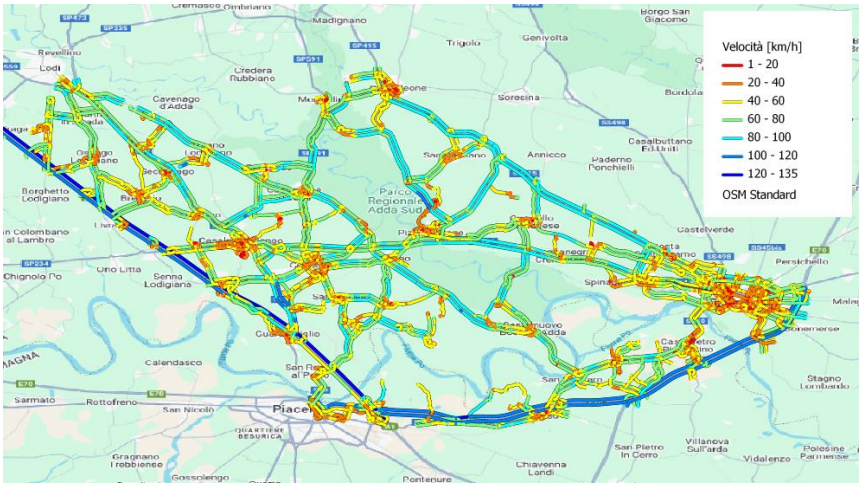

Figure 60: Speed demonstration based on real data from 7 to 8 in the morning

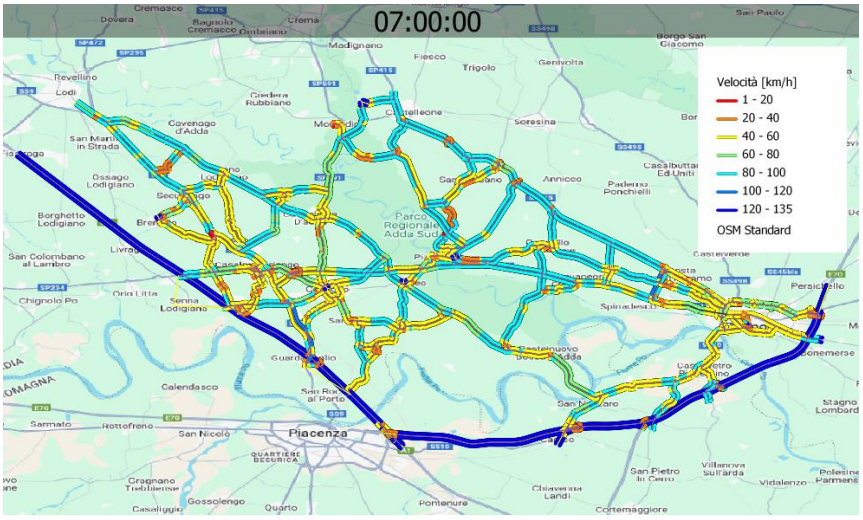

Figure 61: Simulated speed results for a typical Thursday of the day

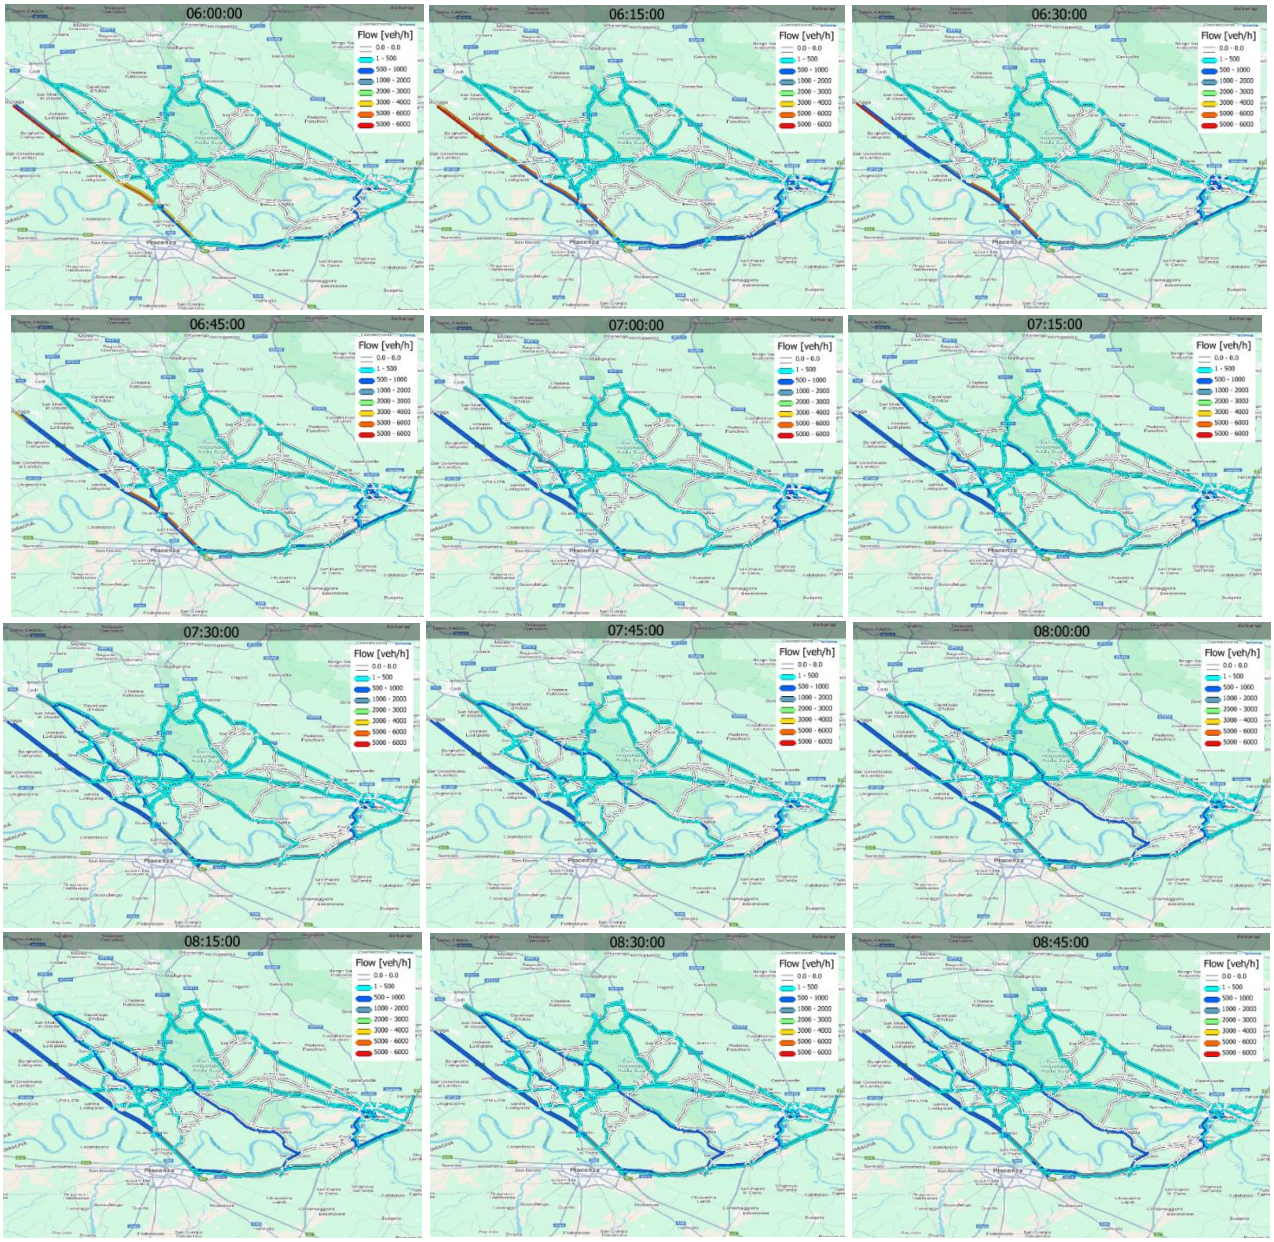

Figure 62. Base scenario simulation O-D Seed

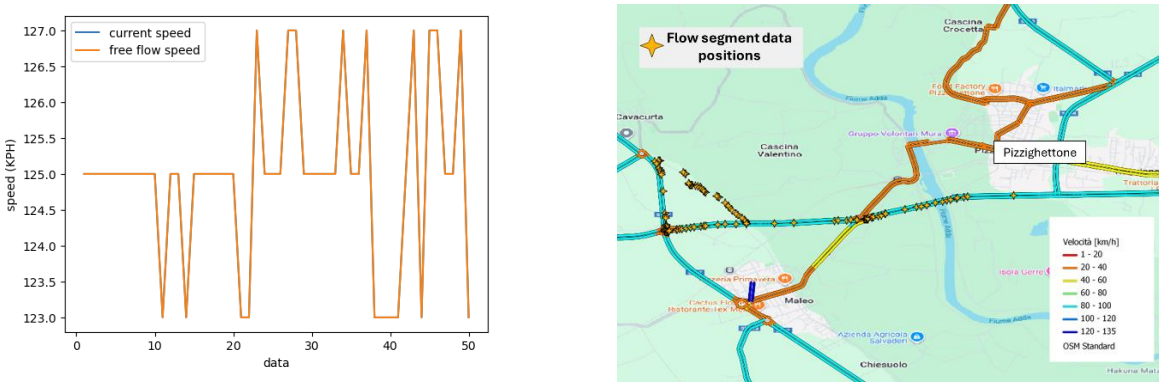

(a) (b)

Figure 63: Comparison between TomTom traffic state data and simulated speeds on road segments approaching the Pizzighettone bridge.

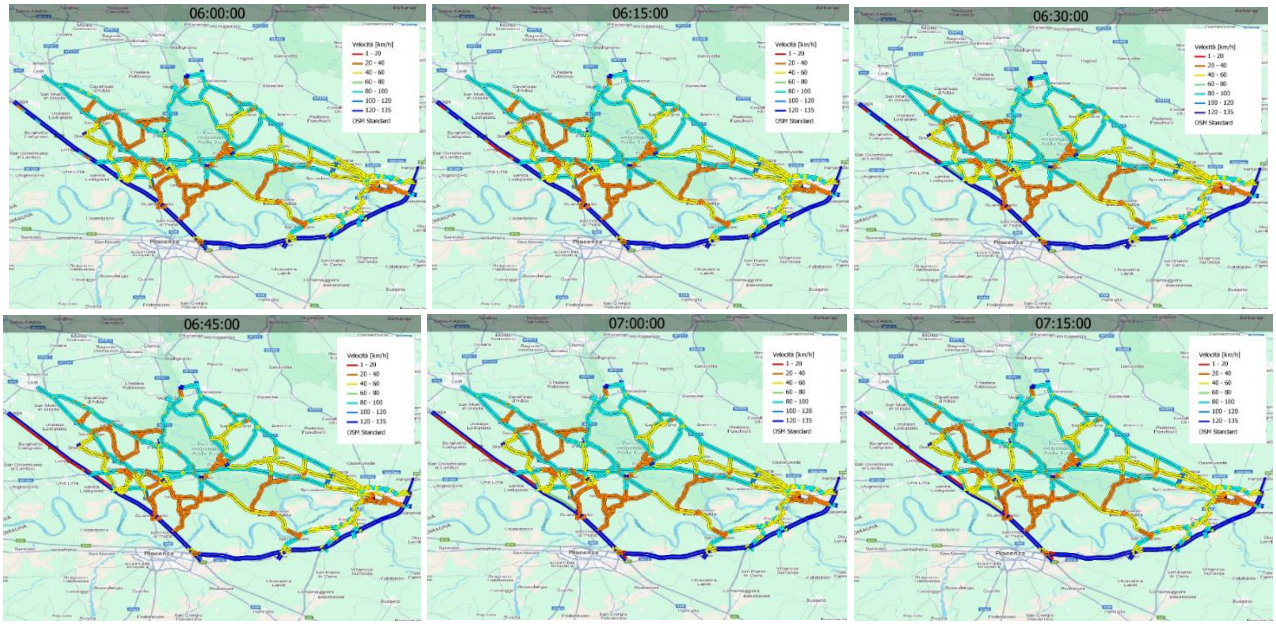

Figure 64: Simulated speed

Scenario A – Bridge closure

Scenario A simulates the closure of a bridge within the study area, representing a high-severity, low-probability event such as one triggered by an earthquake, flooding, or structural degradation. The objective of this scenario is to evaluate the propagation of network-wide impacts resulting from the sudden loss of a major link or node, including traffic redistribution, congestion formation, and increased travel times on alternative routes.

In the simulation model, the bridge and its connecting links are deactivated from 07:15, forcing vehicles to dynamically reroute based on updated travel costs. The model simulates how demand shifts both spatially and temporally, highlighting the most critical alternative corridors and bottlenecks under emergency conditions.

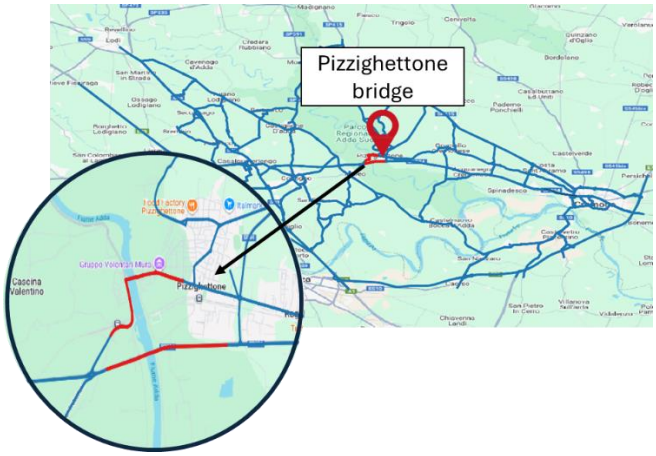

Figure 65: Area affected by closure

During the three-hour simulation period (07:00–10:00), the model accurately captures the impacts of the carriageway closure implemented at 07:00. Once the affected links become unavailable, the routing algorithm successfully reallocates demand by diverting traffic flows to alternative routes. As expected, a significant portion of the diverted traffic is rerouted to Strada Provinciale SP27, which acts as the primary

diversion corridor for maintaining network connectivity and passes through Maleo. A smaller share of the rerouted traffic is instead directed to SP591, which runs near Castiglione d'Adda.

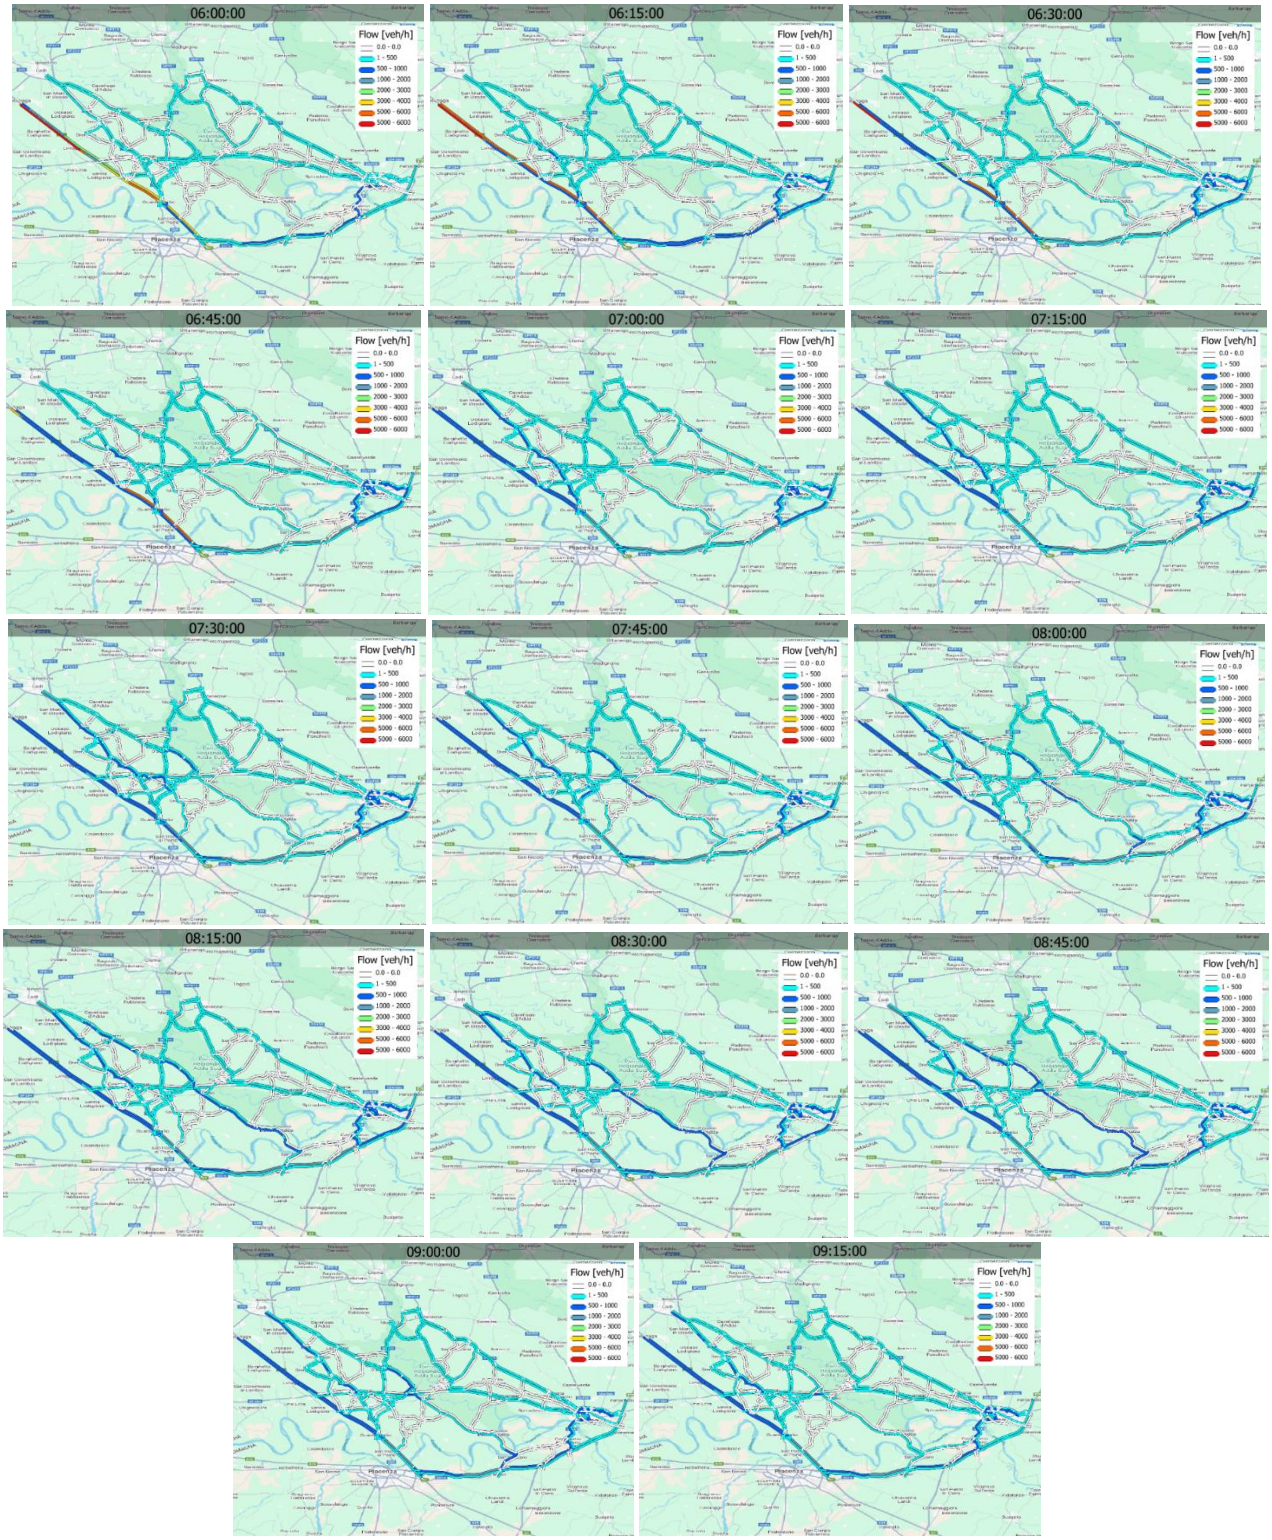

Figure 66: Simulated flow, scenario A-bridge closure

The analysis of simulated speed variations over time (Figure 66) shows that the bridge closure triggers a clear propagation of congestion effects. The reduction in speed originates at the blocked section and progressively propagates upstream along the approach links, forming a backward-moving congestion wave. This behaviour is consistent with well-known traffic phenomena, whereby sudden capacity reductions generate upstream shockwaves.

The central area of the network – particularly the sections approaching the closed bridge – experiences marked speed reductions and partial flow blockage, reflecting queue formation and a local concentration of congestion.

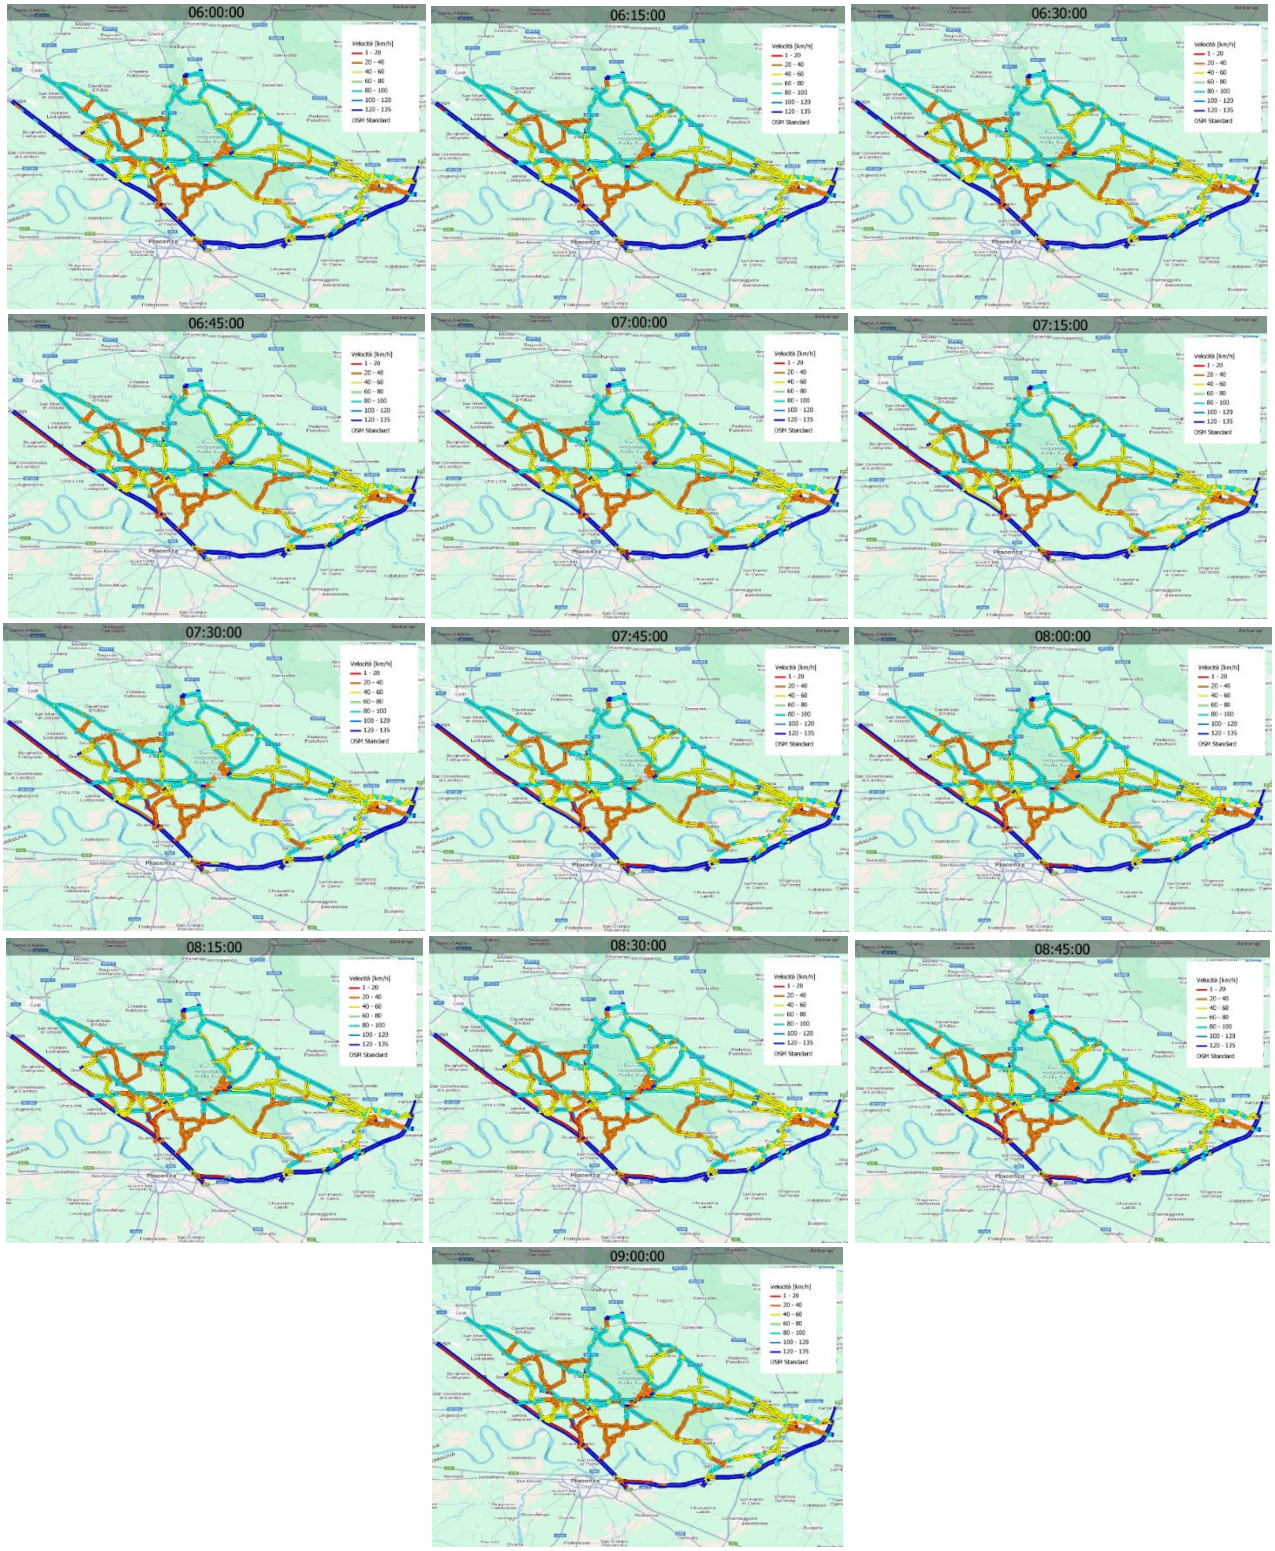

Figure 67. Simulate speed scenario A- Bridge closure

4.3 Findings and Discussion

By enabling a dynamic representation of traffic behaviour under both normal and disrupted conditions, simulation models allow for the quantitative evaluation of network robustness, adaptive capacity, and recovery potential. In the context of natural hazards or infrastructure failures, such tools are essential for anticipating system responses, identifying critical links, and designing mitigation or rerouting strategies that ensure continuity of mobility and access to essential services.

This study contributes to a deeper understanding of how a road network reacts to disruptions and how alternative routes can support mobility continuity during critical events. Through the development and application of a dynamic traffic simulation model, both normal and disrupted operating conditions were analysed, highlighting how traffic flows reorganise and how congestion propagates following the loss of a major link. The results show that the model is able to realistically reproduce rerouting mechanisms and congestion dynamics, providing valuable insights into the network's adaptive behaviour and structural weaknesses.

Beyond the specific case of the Pizzighettone bridge, this work offers a practical foundation for future analyses of transport resilience and recovery planning. It demonstrates how simulation can be used not only to test what-if scenarios, but also to support the definition of meaningful indicators for assessing network resistance and recovery capacity. Overall, the study provides both methodological and analytical contributions to understanding traffic system behaviour under stress and to supporting data-driven decision-making in disaster and resilience management.

Future work should focus on improving model calibration through the use of additional observed data, extending the simulations to multiple time windows and event types, and integrating the framework with real-time monitoring platforms, thereby enhancing its effectiveness as a decision-support tool for resilient traffic management.

5. Conclusions

This deliverable demonstrates that **large-scale simulation-based analyses** provide a robust and effective means to identify potential critical components, systemic vulnerabilities, and cascading effects in complex infrastructure networks. Across water distribution, railway systems, and road traffic, a consistent pattern emerges: **only a limited subset of components or failure scenarios accounts for most system-level impacts**, while the majority of potential failures result in limited or manageable effects.

This observation has clear implications for **risk-informed prioritisation** and confirmed the intrinsic criticalities of the railway infrastructure, already known to the infrastructure manager and for which it has specific procedures, and constitutes a solid basis for the identification of particular cases on which to develop further analyses.A central contribution of this work is the demonstration of the **feasibility and value of exhaustive – or near-exhaustive – simulation campaigns**. While individual simulations rarely reveal critical vulnerabilities, the systematic exploration of thousands of scenarios enables the identification of **non-intuitive weaknesses** that would be extremely difficult, if not impossible, to detect through expert judgement alone.

The analyses also confirm that advanced commercial simulation tools – originally developed for planning or operational purposes – can be effectively repurposed for **resilience-oriented studies**, provided that suitable methodological adaptations and post-processing strategies are applied.

Across all domains, a key insight emerges clearly: *human expertise remains indispensable for interpreting and managing critical situations, but only systematic machine-based simulation can reliably identify which scenarios truly deserve attention*. Simulation should therefore not be seen as a substitute for expert decision-making, but rather as a **necessary preliminary step** that filters, structures, and prioritises expert analysis.

The adoption of an **all-hazards, failure-oriented approach** proved particularly effective. By deliberately decoupling the identification of critical components from specific triggering events, the analyses show that elements identified as critical under generic failure conditions are precisely those that warrant deeper, hazard-specific investigation in subsequent phases. This confirms that **hazard-agnostic screening** is a powerful and efficient first step in resilience assessment, allowing hazard modelling, fragility analysis, and cost–benefit evaluations to be focused where they are most likely to influence system performance.

The results also emphasise the role of **interdependencies**, especially those related to energy supply, control systems, and the spatial co-location of assets. In the water domain, reliance on the power grid emerges as a dominant vulnerability under blackout scenarios. In the transport domains, by contrast, localised failures propagate through network topology, generating congestion and service disruption well beyond the initial point of failure. These outcomes highlight the limitations of asset-by-asset assessments and reinforce the need for **system-level resilience evaluations**, where cascading effects and cross-domain interactions are explicitly addressed.

From an operational standpoint, several consistent lessons can be drawn. **Local disruptions may produce non-local consequences**, particularly in highly interconnected networks. **Redundancy alone is insufficient** unless it is supported by operational feasibility, such as effective rerouting options or emergency power availability. Furthermore, **critical loads and service types** – including hospitals, freight corridors, and emergency services – should play a central role in prioritisation, rather than aggregate demand metrics alone.

In this context, the possibility of linking simulation outputs to **GIS-based representations** significantly enhances their practical value. Such integration enables results to be embedded into existing decision-support systems, facilitating adoption by operators without requiring direct interaction with specialised simulation tools.

Overall, this deliverable shows that **resilience is not achieved by analysing everything in equal detail, but by identifying where focused attention is most needed**. Large-scale, systematic simulation provides the analytical lens required to make this distinction explicit, actionable, and defensible—supporting more informed, proportionate, and effective decisions for the protection of critical infrastructures.

6. References

FEMA (2024). Hazus 6.1: Hazus Earthquake Model, Technical Manual. Federal Emergency Management Agency, Washington DC. https://www.fema.gov/flood-maps/tools-resources/flood-map-products/hazus/user-technical-manuals

FEMA (2024). Hazus 6.1: Hazus Inventory, Technical Manual. Federal Emergency Management Agency, Washington DC.

FEMA (2024). Hazus 6.1: Hazus Flood Model, Technical Manual. Federal Emergency Management Agency, Washington DC.

O' Rourke, T.D., Toprak, S, Sano, Y. (1998). Factors affecting water supply damage caused by the Northridge earthquake. Proceedings of the Sixth US National Conference on Earthquake Engineering.

Hoogendoorn, S. P., Bovy, P. H. L. (2001). State-of-the-art of vehicular traffic flow modelling. Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering, 215(4), 283–303. https://doi.org/10.1177/095965180121500402

van Wageningen-Kessels, F., van Lint, H., Vuik, K., Hoogendoorn, S. P. (2015). Genealogy of traffic flow models. EURO Journal on Transportation and Logistics, 4(4), 445–473. https://doi.org/10.1007/s13676-014-0045-5