Variational Autoencoder-Assisted Accelerated Benders Decomposition for Resilient Stochastic Air Transport Network Reconfiguration under Airspace Disruptions

Iman Rahimi
iman.rahimi-1@uts.edu.au

University of Technology Sydney

Research Article

Keywords: Benders Decomposition, VAE, Deep Learning, Stochastic Optimization

Posted Date: May 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9673076/v1

Additional Declarations: The authors declare no competing interests.

Variational Autoencoder-Assisted Accelerated Benders Decomposition for Resilient Stochastic Air Transport Network Reconfiguration under Airspace Disruptions

Iman Rahimi*

Faculty of Engineering &Information Technoogy, University of Technology Sydney, Australia,
(*corresponding author: iman.rahimi-1@uts.edu.au)

Abstract- This study proposes a Variational Autoencoder (VAE)-assisted stochastic optimization framework for resilient air transport network reconfiguration under large-scale airspace disruptions. The proposed bi-objective model simultaneously minimizes total operational cost and expected unmet demand, enabling analysis of resilience–cost trade-offs under uncertain disruption conditions. To represent disruption uncertainty more realistically, a VAE is employed to generate synthetic disruption scenarios involving airport-capacity reductions, route closures, routing-cost escalation, and demand variation. The generated scenarios are clustered and integrated into a stochastic mixed-integer optimization framework. To solve the resulting large-scale stochastic problem efficiently, an accelerated multi-cut Benders decomposition algorithm incorporating scenario-specific cuts, trust-region stabilization, and cut-management mechanisms is developed. Computational experiments on benchmark-style air transport networks demonstrate that the proposed framework substantially outperforms classical single-cut and classical multi-cut Benders approaches in convergence speed, runtime, and iteration efficiency. The proposed method reduces the number of Benders iterations from more than 80 iterations to fewer than 10 in most tested cases, while reducing generated cuts by up to 90%. In high-penalty disruption scenarios, the classical multi-cut approach required nearly 29,000 seconds of runtime, whereas the proposed framework maintained consistently low computational time across all disruption levels. The results further reveal that reserve capacity becomes concentrated on strategically important hubs, particularly DXB and IST, under severe disruption conditions. Overall, the study highlights the value of integrating deep generative learning with accelerated stochastic decomposition for scalable and resilience-oriented air transport network planning under disruption.

Keywords- Benders Decomposition, VAE, Deep Learning, Stochastic Optimization

1. Introduction

Air transport networks are highly sensitive to geopolitical disruptions, particularly in regions where critical airspace corridors and hub airports are concentrated. Recent escalations in Middle East conflicts (2026)[1] have led to partial airspace closures, increased safety advisories, and widespread rerouting of flights. Major hub airports such as Dubai and Doha, traditionally serving as central connectors between Europe, Asia, and Africa, have experienced operational disruptions, forcing airlines to reconfigure routes and shift traffic toward alternative hubs such as Istanbul and Southeast Asian gateways. These disruptions not only increase operational costs and delays but also expose structural vulnerabilities in global air transport networks.

Despite a growing body of research on airline disruption management and resilient transportation systems, existing studies often focus on short-term operational recovery (e.g., flight delays, crew rescheduling) rather than large-scale network-level reconfiguration under geopolitical uncertainty. Moreover, most traditional optimization models assume static or deterministic disruption patterns, which fail to capture the dynamic and spatially heterogeneous nature of conflict-induced risks. At the same time, machine learning approaches for disruption prediction have largely been applied to weather-related delays or airport congestion, with limited integration into network-level decision-making frameworks.

To address these limitations, this study proposes a data-driven framework for resilient air transport network design under airspace disruption. The framework integrates machine learning–based disruption prediction with a robust optimization model for network reconfiguration. Specifically, a predictive model is developed to estimate disruption risk across air routes based on conflict intensity, airspace restrictions, weather conditions, and fuel price dynamics. These predictions are then embedded into a multi-objective optimization model that determines optimal rerouting strategies, hub reallocation, and capacity adjustments to enhance network resilience while minimizing cost and delay.

The main contributions of this study are threefold. First, it introduces a bi-objective model for reconfiguration for air traffic under disruption considering a trade-off between cost and expected

---

[1] 2026 Iran war - Wikipedia

unmet demand. Second, it uses a VAE for scenario generation using real-datasets. Third, it designed an accelerated multi-cut Benders decomposition. Fourth, it provides empirical insights into the resilience of global air transport networks under Middle East airspace disruptions, highlighting the role of alternative hubs and route diversification strategies.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature on airline disruption management, resilient supply chain design, and machine learning applications in transportation. Section 3 presents the research methodology. Section 4 discusses problem statement and data sources. Section 5 shows the results section. Sections 6 and 7 the discussion and conclusion.

2. Literature Review

Airline disruption management has received substantial attention in aviation operations research, particularly in relation to short-term recovery decisions after unexpected disturbances (Hassan, L.K. et al. 2021; Su Y, et al. 2021). The main streams of this literature include aircraft recovery, crew recovery, passenger reaccommodation, and integrated recovery models. Because the full airline recovery problem is large and highly complex, many studies decompose it into sequential subproblems, where aircraft recovery is usually solved first, followed by crew recovery and passenger recovery. This structure reflects the practical operation of airline control centres, but it also means that many models are designed for tactical or operational recovery rather than strategic network redesign. Recent review studies confirm that most airline disruption management models focus on restoring disrupted schedules, minimizing delay propagation, reducing cancellations, and controlling passenger inconvenience through mixed-integer programming, heuristic, metaheuristic, or decomposition-based methods.

Classical airline disruption models are effective when the disruption is temporary, localized, and operationally manageable, such as aircraft unavailability, airport congestion, weather events, crew misconnection, or short-term capacity reductions. However, conflict-induced disruptions are structurally different. Geopolitical conflicts can lead to airspace closures, route bans, security restrictions, fuel-supply shocks, and sudden hub-capacity changes. These disruptions do not merely delay flights; they can change the feasible topology of the air transport network itself. Therefore, traditional aircraft, crew, and passenger recovery models are not sufficient for analysing large-scale air transport reconfiguration under geopolitical uncertainty. This gap is important because modern air transport networks are highly interconnected, and disruptions at key hubs or airspaces can propagate across regions through rerouting, congestion, and capacity spillover effects. Reviews of air transportation as complex networks also show that connectivity, robustness, delay propagation, and resilience have become major research themes, but many studies remain analytical or descriptive rather than prescriptive decision-support models (Sun, & Wandelt 2021; Gopalakrishnan, & Balakrishnan, 2021)).

The concept of resilience has become central in transportation and supply chain network design. In this context, resilience generally refers to the ability of a system to absorb disruptions, maintain an acceptable level of service, adapt to degraded conditions, and recover after disturbance. Transportation and logistics studies have addressed resilience through redundancy planning, backup capacity, facility fortification, alternative routing, inventory buffers, flexible sourcing, and multi-period recovery decisions. In supply chain network design, stochastic and robust optimization models are widely used to represent disruption risks and uncertain operational conditions. Similar ideas are applicable to air transport networks, where backup hubs, reserve capacity, rerouting options, and scenario-dependent service constraints can be used to improve system resilience. However, many resilience models rely on a limited set of manually specified scenarios, which may not represent the diversity and severity of real-world disruptions (Filho, A.,A., et al. 2026; Müllerklein, & Fontaine, (2025)).

Stochastic programming is particularly relevant for resilient network design because it allows strategic first-stage decisions to be evaluated against multiple disruption scenarios. In two-stage stochastic models, the first stage usually determines preparedness decisions such as facility activation, backup capacity, or network configuration, while the second stage captures scenario-specific operational decisions such as flow routing, unmet demand, and congestion. This structure is well suited to air transport disruption planning because strategic decisions must be made before the exact disruption is known, while routing and service decisions depend on the realized scenario. However, stochastic network design models can become computationally expensive when the number of scenarios, airports, routes, and origin-destination pairs increases. Decomposition algorithms were used widely to solve large-scale optimization problems efficiently (Rahimi, I., 2025; Chuong et al., 2025; Paskov, A. S., 2025). Amongst the decomposition algorithms, Benders decomposition has been widely used in stochastic programming and transportation network design because it separates strategic master decisions from operational subproblems and iteratively improves the approximation of the recourse function through cuts (Crainic, T,G., 2021; Bai, R et al. 2014).

Despite its advantages, classical Benders decomposition may converge slowly when the master problem provides weak lower bounds or when the recourse function is approximated by aggregated cuts. This issue is especially relevant in stochastic network design, where each disruption scenario may have different routing costs, capacity reductions, and congestion patterns. Recent studies have explored enhanced, accelerated, and partial Benders decomposition methods to improve convergence in transportation and logistics problems. Multi-cut Benders decomposition is one important acceleration strategy because it generates separate cuts for each scenario instead of aggregating all scenario information into a single cut. This preserves scenario-specific information and can strengthen the master problem. Other acceleration strategies include stabilization, trust-region restrictions, cut selection, warm-starting, and scenario reduction. These methods are particularly useful when the model contains many scenario-dependent subproblems (Magnanti & Wong (1981); Rahmaniani, R., et al.(2018)).

In parallel, machine learning has increasingly been applied in aviation analytics. A large body of work focuses on flight delay prediction, passenger demand forecasting, delay propagation, and airport congestion prediction. Classical machine learning methods such as random forests, support vector machines, gradient boosting, and neural networks have been used to predict delays from weather, schedule, airport, and historical operational features. More recently, deep learning and graph-based models have been introduced to capture spatial and temporal dependencies in aviation networks. These approaches are motivated by the fact that flight delays and disruptions propagate through network connections rather than occurring independently at each airport. Recent studies on spatial-temporal graph learning and graph machine learning show that airport and route interactions can be represented as dynamic networks, allowing models to learn delay propagation and operational dependencies more effectively than purely tabular models (Kim, Y. J.et al. (2016); Gui, G.et al. (2019)).

However, most machine learning studies in aviation remain prediction-oriented. They estimate delays, demand, congestion, or disruption probability, but they rarely translate these predictions into optimization-based decisions. This creates a methodological gap between predictive analytics and prescriptive optimization. For disruption management, prediction alone is not enough: decision-makers need to know which backup hubs to activate, how much reserve capacity to allocate, which routes to prioritize, and how much unmet demand or congestion is acceptable under different service-level requirements. Therefore, there is a need for predictive-optimization frameworks that convert data-driven disruption information into structured optimization inputs. In this study, this role is assigned to the scenario-generation module: machine learning is used not as a direct decision-maker, but as a mechanism for generating representative disruption scenarios that feed the stochastic optimization model.

Generative machine learning is especially relevant for scenario generation (Dong, et al. (2022; Liang & Tang (2019), Rahimi, I., (2025)). Instead of relying only on manually designed disruption cases, a generative model can learn latent patterns of disruption from scenario features and produce additional plausible scenarios . Variational autoencoders are suitable for this purpose because they learn a low-dimensional latent representation of complex data and generate new samples by decoding points from the latent space (Carbonera, M.et al.; (2024)). In the context of air transport disruption, scenario features can include airport capacity reductions, route availability, cost inflation, demand variation, and disruption severity. Once generated, these scenarios can be filtered, clipped to operationally meaningful ranges, and clustered to obtain a smaller set of representative scenarios. Scenario clustering is important because solving the stochastic optimization model over all generated scenarios would be computationally expensive. By using cluster representatives and assigning probabilities according to cluster size, the model can preserve key disruption patterns while keeping the optimization tractable.

Based on the above literature, three major research gaps can be identified. First, most airline disruption management studies focus on operational recovery and do not adequately address large-scale strategic network reconfiguration under geopolitical uncertainty. Second, resilience-oriented transportation models often rely on static, manually defined, or limited disruption scenarios, which restricts their ability to represent dynamic and evolving risk conditions. Third, machine learning models in aviation are mainly used for prediction and are rarely integrated into stochastic optimization frameworks for decision support. This study addresses these gaps by proposing a bi-objective stochastic optimization framework for resilient air transport network design under conflict-induced disruptions. The framework combines VAE-based disruption scenario generation, scenario clustering, and an accelerated multi-cut Benders decomposition algorithm. The model explicitly captures the trade-off between total system cost and expected unmet demand, while the decomposition structure separates strategic backup-capacity decisions from scenario-specific routing, congestion, and service-level decisions.

3. Research Methodology

This study proposes a VAE-assisted stochastic optimization framework for resilient air transport network reconfiguration under airspace disruptions. The proposed methodology integrates benchmark-inspired air transport network data, deep learning-based disruption scenario generation, stochastic network optimization, and accelerated decomposition techniques to support resilience-oriented decision making under uncertainty. The overall framework consists of four main stages. The input data were compiled from publicly available sources, including OpenSky Network air-traffic data [2], Natural Earth GIS base layers [3], and the World Port Index [4] for port locations and characteristics. This data were further processed to construct the transport network, disruption scenarios, demand parameters, and service-level constraints.

In Stage 1, a scenario-based air transport network dataset is constructed and preprocessed using a benchmark-style network structure containing hub information, origin–destination (OD) demand pairs, path availability, routing costs, and hub capacity data under multiple disruption scenarios. The dataset is transformed into an optimization-ready format including disrupted route availability matrices, scenario-dependent demand realizations, hub capacities, and reserve allocation parameters. The resulting network representation forms a directed stochastic transportation structure suitable for resilient reconfiguration analysis under uncertain disruption conditions.

[2] https://opensky-network.org/data
[3] https://www.arcgis.com/home/item.html
[4] https://www.naturalearthdata.com

In Stage 2, disruption characteristics are modeled through stochastic variations in airport capacities, route availability, routing costs, and OD demand patterns. To enhance scenario diversity and capture nonlinear disruption behaviors, a Variational Autoencoder (VAE) is trained on the generated disruption feature space. The VAE learns latent representations of disruption patterns and subsequently generates additional synthetic disruption scenarios reflecting complex operational uncertainties. The generated scenarios are postprocessed using feasibility filtering and normalization procedures to ensure operational consistency and realistic network behavior. Representative scenarios and associated probabilities are then used to construct the stochastic optimization environment.

Stage 3 formulates a bi-objective stochastic air transport network reconfiguration model that simultaneously minimizes total operational cost and expected unmet demand under disruption scenarios. The model determines strategic hub activation decisions, reserve capacity allocation, and disruption-aware traffic rerouting decisions across the network. The conflicting objectives of operational efficiency and resilience are balanced using a penalty-based scalarization approach controlled by the unmet-demand penalty parameter (λ). Different λ values are evaluated to investigate the trade-off between service continuity and operational expenditure under varying disruption conditions.

Finally, Stage 4 implements an accelerated multi-cut Benders decomposition framework to solve the large-scale stochastic optimization problem efficiently. The proposed solution approach incorporates scenario-specific multi-cut generation, trust-region stabilization, and cut-management mechanisms to improve convergence stability and computational scalability. The proposed accelerated approach is compared with classical single-cut and classical multi-cut Benders decomposition methods using several performance metrics, including total objective value, expected unmet demand, runtime, iteration count, number of generated cuts, lower and upper bound convergence behavior, and optimality gap. In addition to algorithmic evaluation, the resulting solutions are analyzed from a resilience and network-management perspective through cost–service trade-off analysis, reserve allocation behavior, hub utilization patterns, and critical hub identification under airspace disruptions (Figure 1).

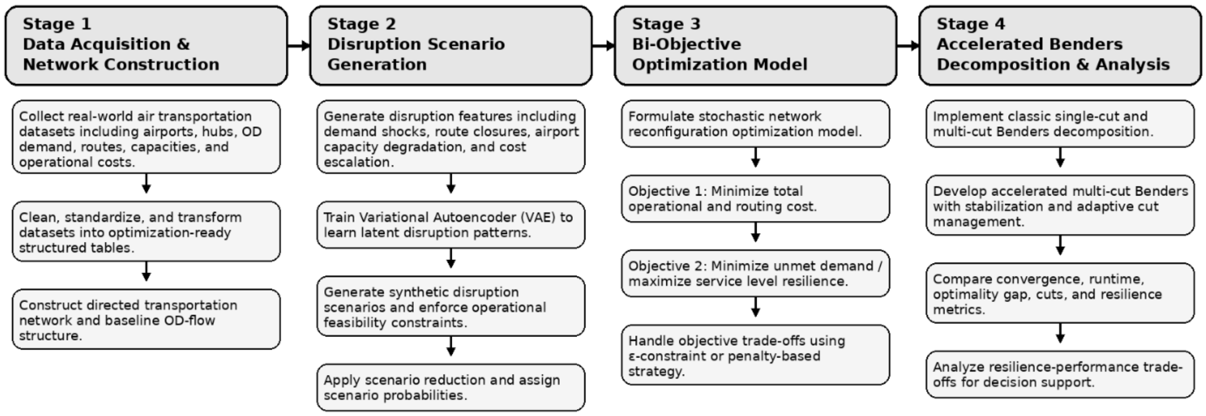

Figure 1 Research Methodology Flowchart

3.1 Scenario construction

To represent uncertainty in the resilient air transport network, a VAE-assisted disruption scenario generation framework is developed. First, a benchmark-style scenario-based air transport dataset is constructed using stochastic variations in airport capacities, route availability, routing costs, and origin–destination (OD) demand patterns under airspace disruption conditions. Each disruption scenario is represented through a compact feature structure capturing operational changes in airport capacity, disrupted route accessibility, traffic-flow cost escalation, and demand fluctuation behavior. These features are used to characterize heterogeneous disruption conditions across the transportation network.

A Variational Autoencoder (VAE) is then trained to learn latent representations of the disruption patterns embedded within the scenario feature space. The encoder maps the disruption features into a low-dimensional latent space, while the decoder reconstructs representative disruption structures from sampled latent vectors. After training, additional synthetic disruption scenarios are generated by sampling from the latent space and decoding the samples into new disruption realizations. To preserve operational realism and avoid infeasible network states, the generated scenario attributes are clipped and normalized within predefined operational bounds. This learning-assisted generation process enables the framework to capture nonlinear disruption relationships and produce diverse disruption conditions beyond the initially constructed scenarios.

Since directly solving the stochastic optimization model over all generated scenarios would significantly increase computational complexity, a scenario reduction stage is incorporated. The generated disruption scenarios are clustered using K-means clustering, and representative scenarios are selected from the resulting clusters. The relative size of each cluster is used to estimate the corresponding scenario probability. The clustered scenarios are then transformed into optimization-ready stochastic inputs, including scenario-dependent airport capacities, path availability matrices, routing costs, OD demand realizations, and scenario probabilities. This modular scenario-generation structure separates uncertainty modeling from optimization and allows the same optimization framework to be evaluated under manually constructed, VAE-generated, or clustered disruption environments.

The resilient air transport network reconfiguration problem is formulated as a bi-objective stochastic optimization model. The model simultaneously minimizes total operational cost and expected unmet demand under disruption scenarios. Strategic decisions include hub activation and reserve-capacity allocation, while operational decisions determine traffic rerouting across the disrupted network. The trade-off between resilience and operational efficiency is controlled through the unmet-demand penalty parameter λ, allowing different resilience preferences to be evaluated systematically.

To solve the resulting large-scale stochastic optimization problem efficiently, an accelerated multi-cut Benders decomposition framework is implemented. The decomposition separates strategic reserve-allocation decisions from scenario-specific operational routing decisions. The master problem determines hub activation and reserve-capacity allocation decisions represented by binary variables $y_h$and continuous reserve variables q_h, respectively. Scenario-specific recourse approximation variables $\theta_s$are also included in the master problem. For fixed reserve-capacity decisions, each scenario subproblem solves a linear traffic-routing problem under disrupted network conditions, determining OD flows and unmet demand across the network.

The dual values associated with the scenario-specific capacity constraints are used to generate Benders optimality cuts that are iteratively added to the master problem. A multi-cut strategy is adopted in which each disruption scenario contributes an independent optimality cut during every iteration. Compared with classical single-cut aggregation, this approach preserves scenario-specific information and improves convergence behavior. To further improve computational performance, three acceleration mechanisms are incorporated. First, the VAE-generated disruption scenarios are clustered before optimization to reduce the number of stochastic subproblems while maintaining representative disruption diversity. Second, the multi-cut strategy strengthens the approximation of the stochastic recourse function by preserving scenario-level information. Third, a trust-region stabilization mechanism is applied to the reserve-capacity decisions in the master problem to prevent excessive oscillation between consecutive iterations and improve convergence stability.

The resulting framework forms a VAE-assisted accelerated multi-cut Benders decomposition approach for resilient air transport network reconfiguration under airspace disruptions. The main methodological contribution of the study is the integration of deep learning-based disruption scenario generation, scenario clustering, and accelerated stochastic decomposition within a unified resilience-oriented decision-support framework for disrupted air transport systems.

4. Problem Statement

The global air transport network is a highly interconnected and hub-dependent system that supports large-scale passenger mobility, cargo transportation, and international economic connectivity. Modern airline operations rely heavily on hub-and-spoke structures in which a limited number of strategically positioned airports facilitate long-haul connectivity and traffic redistribution across continents. While this structure improves operational efficiency under normal conditions, it simultaneously increases the vulnerability of the network to large-scale disruptions affecting critical hubs or major air corridors.

Recent airspace disruptions and geopolitical instability in the Middle East have demonstrated the fragility of global aviation systems under rapidly changing operational conditions. Airspace restrictions, partial route closures, increased safety concerns, and regional conflict escalation have forced airlines to reroute flights, modify schedules, and redistribute traffic toward alternative hubs. As a result, several major transfer airports have experienced sudden changes in traffic flow, congestion levels, and operational utilization. These disruption-driven changes significantly affect network efficiency, operational cost, and service continuity across interconnected aviation systems.

Traditional airline network planning and optimization approaches are generally developed under assumptions of relatively stable operational conditions. Many existing models focus primarily on deterministic cost minimization and are not designed to adapt dynamically to large-scale uncertain disruptions. Furthermore, most conventional network optimization frameworks do not explicitly incorporate disruption-driven uncertainty generation, resilience-oriented reserve allocation, or adaptive network reconfiguration under severe airspace constraints. As a result, these approaches often fail to provide effective decision-support capabilities during rapidly evolving disruption conditions.

Air transport systems also operate as complex multi-stakeholder infrastructures involving airlines, airport operators, regulatory authorities, cargo operators, and passengers with partially conflicting objectives. Airlines aim to maintain operational efficiency and minimize rerouting costs while preserving schedule reliability and service continuity. Airport operators face congestion risk and limited processing capacity, particularly when traffic is redistributed toward alternative hubs during disruption events. Regulatory authorities impose safety-related operational constraints through route restrictions and airspace-control policies, while passengers and cargo stakeholders require reliable transportation and minimal service interruption. These competing objectives become increasingly difficult to balance under severe airspace disruptions characterized by uncertain route availability, fluctuating demand, and dynamically changing operational conditions.

The disruption propagation behavior of hub-and-spoke aviation systems further increases the complexity of resilient network management. Disruptions affecting a small number of strategically important hubs may propagate across the global network through rerouted traffic flows, congestion amplification, and connectivity loss. Consequently, resilience-oriented air transport planning requires not only disruption-aware operational routing decisions, but also strategic reserve-capacity allocation and adaptive network reconfiguration mechanisms capable of preserving service continuity under uncertain disruption scenarios.

To address these challenges, this study proposes a VAE-assisted stochastic optimization framework for resilient air transport network reconfiguration under airspace disruptions. The proposed framework integrates data-driven disruption scenario generation with stochastic network optimization and accelerated decomposition techniques. A Variational Autoencoder (VAE) is employed to learn latent disruption patterns and generate synthetic but operationally realistic disruption scenarios involving airport-capacity reduction, route closure, routing-cost escalation, and OD-demand variation. These disruption scenarios are then integrated into a stochastic resilient network reconfiguration model that simultaneously minimizes total operational cost and expected unmet demand. The resulting problem is characterized by several major challenges including disruption uncertainty and scenario diversity, network interdependencies and disruption propagation, resilience–cost trade-offs, large-scale computational complexity, and adaptive reserve allocation and traffic redistribution.

To manage these challenges efficiently, the proposed framework employs an accelerated multi-cut Benders decomposition algorithm that separates strategic reserve-allocation decisions from scenario-specific operational routing decisions. The overall objective is to provide a scalable and resilience-oriented decision-support framework capable of improving adaptive air transport network performance under uncertain airspace disruptions.

Indexes and sets

P_k	Set of candidate paths for OD pair k	
S	set of disruption scenarios	s∈S
H∈ N	Set of candidate backup hubs	h∈ H
O(k)	origin airport of OD pair k	

Parameters

c_{kps}	routing cost of path p∈P_k in scenario s
F_h	fixed cost of activating backup hub h
C_h^R	unit reserve-capacity cost of hub h

B	Hub budget
a_{kph}	hub-use coefficient: whether/how much path p of OD k uses hub h
λ	Penalty cost for unmet demad
σ_{kps}	1 if path p for OD k is available in scenario s, otherwise 0
d_{ks}	demand of OD pair k in scenario s
u_{hs}	Available capacity of hub h in scenario s
Q_h	Maximum reserve capacity of hub h
P_s	Probability of scenario s

Decision Variables

$y_h \in \{0,1\}$	Backup Activation of hub
$q_h \geq 0$	Reserve capacity assigned to hub h
$x_{kps} \geq 0$	routed flow on path $p$for OD pair $k$under scenario s
$m_{ks} \geq 0$	Unmet demand

Aggerated Objectives: Minimizing total cost and unmet demand

$$\text{Min } Z_1 = \sum_{h\in H} F_h y_h + \sum_{h\in H} C_h^R q_h + \sum_{s\in S} p_s [\sum_{k\in K} \sum_{p\in P} c_{kps} \times x_{kps} + \lambda \times \sum_{k\in K} m_{ks} \tag{1}$$

Equation 1 is weighted objective: minimizing first-stage hub activation and reserve-capacity costs, expected routing costs, and λ-weighted unmet demand penalties across all disruption scenarios.

Constraints:

Equation 2 is flow-balance constraint ensuring that routed flow and unmet demand together satisfy the total OD demand for each OD pair under every scenario.

$$\sum_{p\in P_k} x_{kps} + m_{ks} = d_{ks} \ \forall k \in K, s \in S \tag{2}$$

Equation 3 is hub-capacity constraint limiting total routed flow through each hub to the available scenario-dependent capacity plus allocated reserve capacity.

$$\sum_{k\in K} \sum_{p\in P_k} a_{kph \times x_{kps}} \leq u_{hs} + q_h \ \ \forall h \in H, s \in S \tag{3}$$

Equation 4 is reserve-activation constraint ensuring reserve capacity can only be assigned to activated backup hubs.

$$q_h \leq Q_h \times y_h \quad \forall h \in H \tag{4}$$

Equation 5 is hub-budget constraint limiting the total number of activated backup hubs.

$$\sum_{h \in H} y_h \leq B \tag{5}$$

Equation 6 is route-availability constraint forcing flow on unavailable disrupted paths to zero.

$$x_{kps} = 0 \tag{6}$$

Equation 7 is variable-domain constraints defining binary hub-activation variables and nonnegative continuous decision variables.

$$y_h \in \{0,1\}, q_h \geq 0, x_{kps} \geq 0, m_{ks} \geq 0 \tag{7}$$

Benders decomposition implementation:

The first-stage variables are:

y_h, q_h

The scenario subproblem variables are:

x_{kps}, m_{ks}

Scenario subproblem:

For each scenario s, the following equations are solved:

Equation (8) is scenario-specific subproblem objective minimizing routing costs and unmet-demand penalties for a given reserve-capacity allocation.

$$Q_s(q) = \min \sum_{k \in K} \sum_{p \in P_k} c_{kps} \times x_{kps} + \lambda \sum_{k \in K} m_{ks} \tag{8}$$

Subject to:

Equation (9) is subproblem demand-satisfaction constraint ensuring each OD demand is either served or counted as unmet demand.

$$\sum_{p \in P_k} x_{kps} + m_{ks} = d_{ks} \quad \forall k \in K, s \in S \tag{9}$$

Equation (10) is subproblem hub-capacity constraint restricting scenario flow according to available and reserved hub capacities.

$$\sum_{k \in K} \sum_{p \in P_k} a_{kph \times x_{kps}} \leq u_{hs} + q_h \ \ \forall h \in H, s \in S \tag{10}$$

Equation (11) subproblem route-feasibility constraint prohibiting flow assignment on unavailable paths.

$$x_{kps} = 0 \ \ \forall k, p, with\ \partial_{kps} = 0 \tag{11}$$

$x_{kps} \geq 0, m_{ks} \geq 0$

Let π_{hs} be the dual value of the hub-capacity constraint in scenarios. Then the Benders cut for scenario s is as follows. Equation (12) is multi-cut Benders optimality cut approximating the recourse function for each disruption scenario.

$$\theta_s \geq \alpha_s + \sum_{h \in H} \beta_{hs} q_h \tag{12}$$

Where:

Equation (13) is definition of the Benders cut slope coefficient obtained from the dual variable of the hub-capacity constraint.

$$\beta_{hs} = \pi_{hs} \tag{13}$$

Equation (14) is definition of the Benders cut intercept parameter for each scenario.

$$\alpha_s = Q_s(q) - \sum_{h \in H} \beta_{hs} q_h \tag{14}$$

Multi-cut master problem:

Equation (15) is multi-cut master-problem objective minimizing first-stage costs and expected recourse costs over all scenarios.

$$\text{Min} \sum_{h\in H} F_h y_h + \sum_{h\in H} C_h^R q_h + \sum_{s\in S} \rho_s \theta_s \quad (15)$$

Subject to:

Equation (16) is master-problem hub-budget constraint limiting the number of activated hubs. Equation (17) is master-problem reserve-capacity activation constraint linking reserve allocation to hub activation decisions. Equation (18) is scenario-specific multi-cut Benders optimality constraints added iteratively to approximate recourse costs. Equation (19) is domain constraints for master-problem decision variables.

$$\sum_{h\in H} y_h \le B \quad (16)$$

$$q_h \le Q_h \times y_h \; \forall h \in H \quad (17)$$

$$\theta_s \ge \alpha_{rs} + \sum_{h\in H} \beta_{hs}^r \times q_h \; \forall s \in S, r \in R_s \quad (18)$$

$$y_h \in \{0,1\}, q_h \ge 0, \theta_s \ge 0 \quad (19)$$

Single-cut master problem:

The single-cut version aggregates the scenario cuts into one expected recourse cut:

$$\theta \ge \alpha^r + \sum_{h\in H} \beta_{hs}^r \times q_h \quad (20)$$

Where

Equations (21) and (22) are weighted aggregation of scenario intercept coefficients in the single-cut formulation and weighted aggregation of scenario slope coefficients in the single-cut formulation, respectively.

$$\alpha^r = \sum_{s\in S} \rho_s \times \alpha_{rs} \quad (21)$$

$$\beta_h^r = \sum_{s\in S} \rho_s \times \beta_{hs}^r \quad (22)$$

The single cut master is minimizing first-stage costs and aggregated expected recourse cost (equation (23)).

$$\text{Min} \; \sum_{h\in H} (F_h y_h + \sum_{h\in H} C_h^R q_h) + \theta \quad (23)$$

Subject to:

Equation (24) is single-cut master-problem hub-budget constraint. Equation (25) is single-cut reserve-capacity activation constraint. Equation (26) is aggregated Benders optimality cut for the single-cut master problem. Equation (27) is variable-domain constraints for the single-cut master problem.

$$\sum_{h \in H} y_h \leq B \tag{24}$$

$$q_h \leq Q_h \times y_h \ \forall h \in H \tag{25}$$

$$\theta \geq \alpha^r + \textstyle\sum_{h \in H} \beta_{hs}^r \times q_h \tag{26}$$

$$y_h \in \{0,1\}, q_h \geq 0, \theta_s \geq 0 \tag{27}$$

Accelerated multi-cut master:

The trust-region constraints used in the accelerated master are as follows. Equation (28) is upper trust-region stabilization constraint restricting excessive increases in reserve-capacity decisions between iterations. Equation (29) is lower trust-region stabilization constraint restricting excessive decreases in reserve-capacity decisions between iterations.

$$q_h \leq q_h^c + \Delta \times Q_h \ \forall h \in H \tag{28}$$

$$q_h \geq \max\{0, q_h^c - \Delta \times Q_h\} \tag{29}$$

Where q_h^c is the current trust-region center and Δ is the trust-region radius. Equation (30) is accelerated multi-cut master-problem objective integrating trust-region stabilization with scenario-specific Benders cuts for improved convergence.

$$\text{Min } \textstyle\sum_{h \in H}(F_h y_h + \sum_{h \in H} C_h^R q_h) + \sum_{s \in S} \rho_s \times \theta_s \tag{30}$$

subject to the hub-budget, reserve-activation, multi-cut constraints, and trust-region constraints above. The proposed accelerated multi-cut Benders decomposition is presented in Figure 2.

Step	Operation
1	**Input:** cleaned air transport network data, OD demands, route costs, capacities, candidate hubs, disruption features, penalty grid Λ, tolerance ϵ, maximum iteration limit $T_{\max}$, stabilization parameters, and cut-pool limits.
2	**Generate disruption scenarios using VAE:** train the VAE on disruption-related features extracted from the real data; sample synthetic scenarios from the latent space; clip generated demands, capacities, costs, and route disruptions to feasible operational ranges.
3	Build the final scenario set Ω and assign probabilities p_s. No K-means clustering step is used in this algorithm unless it is explicitly activated in a separate preprocessing script.
4	**For each** penalty value $\lambda \in \Lambda$ **do**
5	Initialize cut pools $\mathcal{C}_s \leftarrow \emptyset$ for all $s \in \Omega$, bounds $LB \leftarrow -\infty$, $UB \leftarrow +\infty$, incumbent solution $x^* \leftarrow \emptyset$, and stabilization center $\bar{x}^1$.
6	**For** iteration $t = 1, 2, \ldots, T_{\max}$ **do**
7	Solve the stabilized multi-cut master problem in Eqs. (1)-(4) and obtain (x^t, η^t).
8	Compute the master lower bound using the original master value: $LB^t = C^{\text{res}}(x^t) + \sum_{s\in\Omega} p_s \eta_s^t$.
9	**For each** scenario $s \in \Omega$ **do**
10	Solve the scenario subproblem with fixed $x = x^t$ and penalty λ.
11	Obtain the recourse value $\theta_s(x^t)$ and the dual multipliers required for the Benders cut.
12	Generate one scenario-specific optimality cut $\eta_s \geq \alpha_s^t + (\beta_s^t)^\top x$.
13	Add the new cut to the scenario cut pool $\mathcal{C}_s$.
14	**End for**
15	Compute the feasible upper bound: $UB^t = C^{\text{res}}(x^t) + \sum_{s\in\Omega} p_s \theta_s(x^t)$.
16	If $UB^t < UB$, update $UB \leftarrow UB^t$ and $x^* \leftarrow x^t$.
17	Update $LB \leftarrow \max\{LB, LB^t\}$ and compute the relative gap: $gap = (UB - LB)/\max\{1, \lvert UB\rvert\}$.
18	Apply cut management: remove inactive, weak, or old cuts while keeping recent active cuts and cuts supporting the current lower bound.
19	Update the stabilization center $\bar{x}^{t+1}$ using the incumbent solution or the accepted master solution.
20	**If** $gap \leq \epsilon$ **then break**.
21	**End for**
22	Store the solution for λ: objective value, true cost without unmet-demand penalty, expected unmet demand, served-demand ratio, runtime, number of iterations, number of cuts, final gap, and convergence history.
23	**End for**
24	**Output:** cost-service trade-off results, Benders convergence plots, runtime comparison, cut-count comparison, and final summary table across all tested λ values and algorithms.

Figure 2 VAE-Assisted Accelerated Multi-Cut Benders Decomposition

5. Results

To validate the proposed bi-objective disruption management framework, the extensive-form mixed-integer model was first solved. Figure 3 (a) illustrates the cost–service trade-off obtained from the proposed resilient air transport network reconfiguration model under different values of the penalty parameter λ. The horizontal axis represents the expected unmet demand, while the vertical axis denotes the operational cost of the network. The results demonstrate a clear trade-off between economic efficiency and service reliability. When smaller values of $\lambda$are used, the optimization model prioritizes operational cost minimization, resulting in lower operating costs but higher levels of unmet demand. As λincreases, the model increasingly penalizes unmet demand, forcing the network to preserve more traffic and improve service continuity. Consequently, the expected unmet demand decreases, although this improvement requires higher reserve utilization and increased operational effort. The nonlinear shape of the curve further indicates diminishing marginal benefits, meaning that beyond a certain resilience level, additional reductions in unmet demand require disproportionately larger operational investments. This behavior confirms that the proposed model successfully captures the resilience–cost balance inherent in real-world transportation systems.

Figure 3 (b) presents the reserve allocation heatmap across different hubs and λvalues. The figure shows how the optimization model strategically allocates reserve capacity to selected airports as resilience importance increases. For low values of λ, almost no reserve capacity is allocated because the model mainly focuses on reducing operational cost. However, as λincreases, reserve resources begin concentrating on strategically important hubs such as DXB, IST, AUH, and CAI. In particular, DXB and IST receive the largest reserve allocations under higher resilience settings, indicating their critical role in maintaining connectivity and absorbing disrupted traffic flows. Rather than distributing reserve uniformly across all airports, the model concentrates backup capacity on a limited number of highly connected hubs, which reflects realistic hub-and-spoke operational strategies commonly observed in international aviation systems. This result suggests that network resilience can be enhanced more effectively through targeted reinforcement of critical hubs rather than through equal reserve distribution across the entire network.

Figure 3 (c) illustrates the hub utilization heatmap under varying λvalues. The color intensity represents the utilization ratio of each hub, where brighter colors indicate higher utilization levels. Under small λvalues, many hubs remain underutilized because the optimization model allows higher unmet demand in order to minimize cost. As the penalty parameter increases, the utilization of major hubs rises significantly, indicating that the network increasingly relies on available airport capacity to preserve service continuity. Several critical hubs approach full utilization under large λvalues, particularly DXB, IST, and CAI, demonstrating their dominant role in rerouting and traffic recovery operations during disruptions. This behavior indicates that the proposed model adaptively restructures traffic flows toward strategically important hubs when resilience becomes a priority.

The utilization heatmap also provides important managerial insights regarding network saturation and operational vulnerability. Although higher λvalues improve demand satisfaction and reduce unmet traffic, they simultaneously push several hubs toward near-capacity operating conditions. This implies that excessive concentration of rerouted traffic on a small number of hubs may increase congestion risk and reduce operational flexibility under severe disruption scenarios. Therefore, the results highlight an important resilience trade-off: improving service continuity often requires intensive utilization of major hubs, which may expose the network to secondary congestion and cascading operational failures. Overall, the figures demonstrate that the proposed optimization framework produces meaningful and interpretable resilience patterns consistent with real-world air transport network behavior.

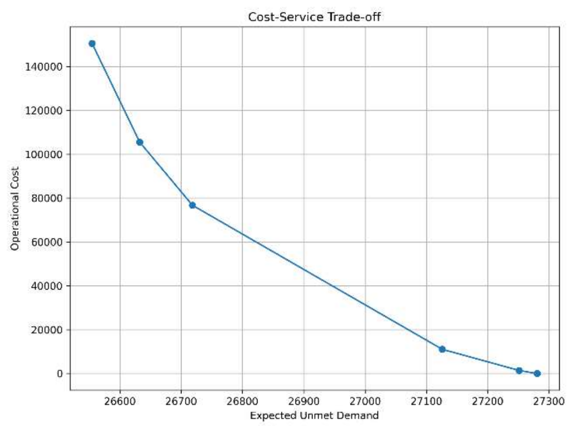

(a) Cost-service trade-off

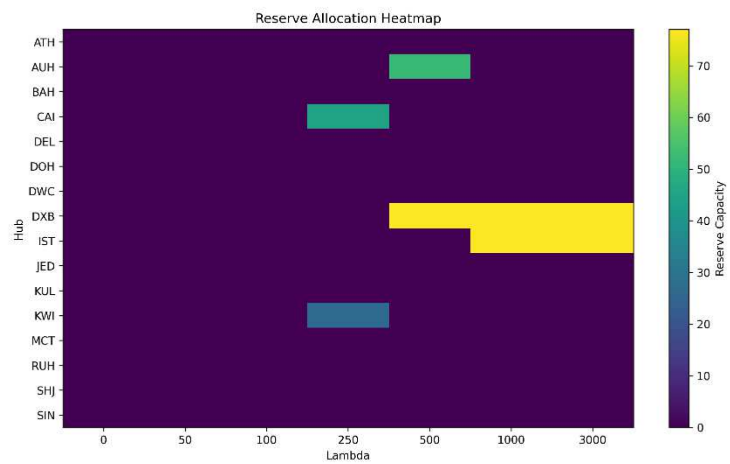

(b)Reserve Allocation Heatmap

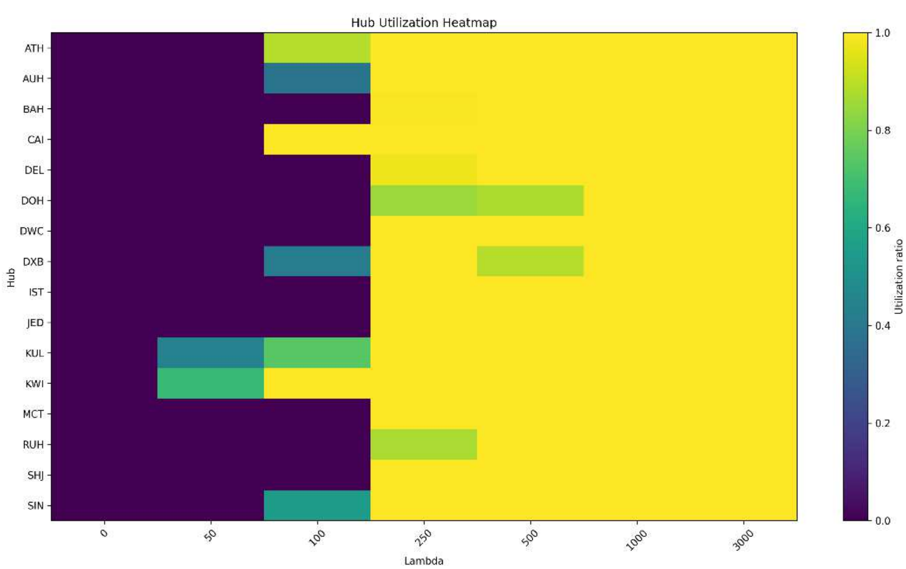

(c) hub utilization heatmap

Figure 3 Hub Utilization Heatmap

From Figures 4-5, the accelerated multi-cut strategy significantly reduces the number of generated cuts compared to the classical multi-cut approach (up to 90% reduction), while maintaining a richer information structure than the single-cut method (through less iterations). This balance leads to a more efficient master problem and contributes to faster convergence.

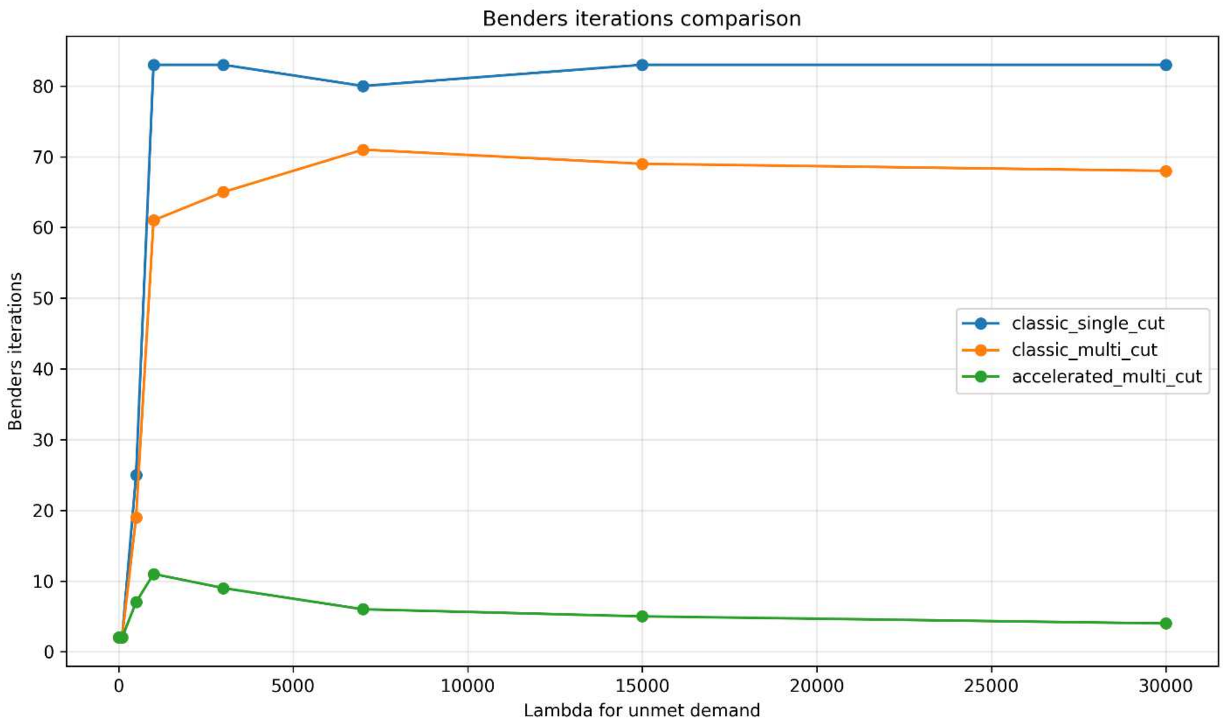

Figure 4 Benders iterations comparison

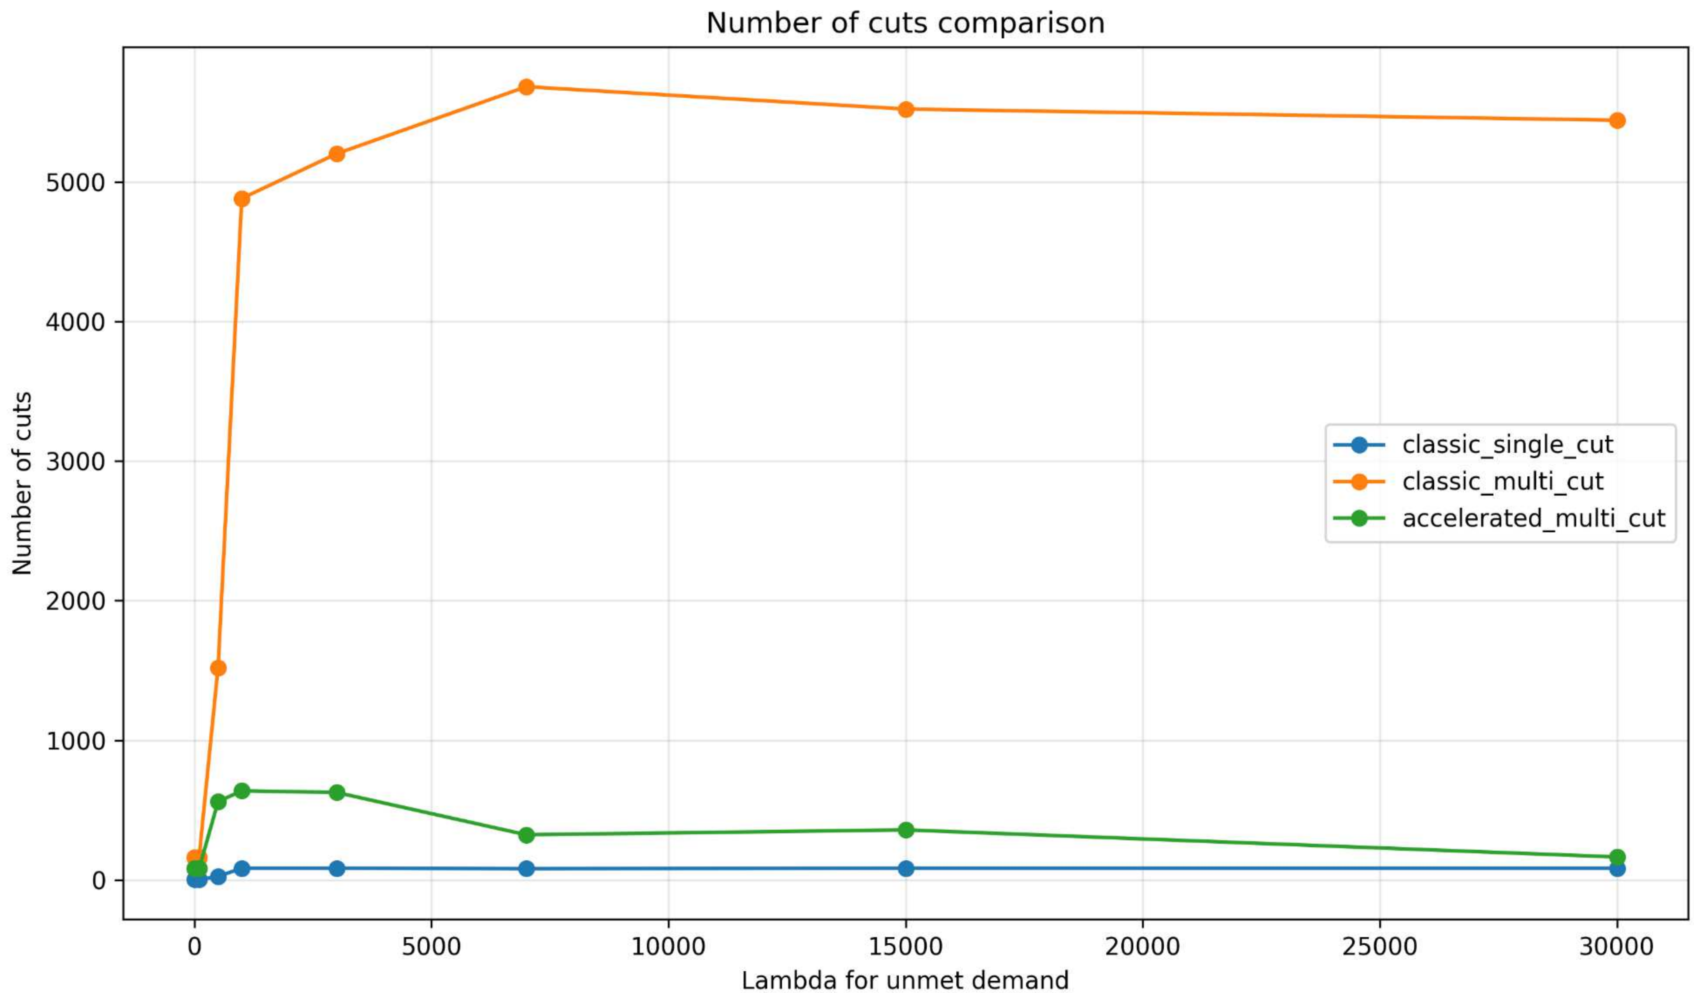

Figure 5 Number of cuts comparison

Figure 6 presents the runtime comparison of the three Benders decomposition strategies—classic single-cut, classic multi-cut, and the proposed accelerated multi-cut—across different values of the unmet demand penalty parameter (λ). The results reveal a clear and consistent performance advantage for the proposed approach. The accelerated multi-cut method maintains a remarkably low and stable runtime across all λ values, demonstrating strong robustness and scalability. In contrast, the classic multi-cut approach exhibits highly unstable behavior, with runtime increasing dramatically for moderate λ values (e.g., peaking at around 29,000 seconds for λ = 7,000). This sharp increase indicates severe computational inefficiency caused by excessive cut generation and the resulting complexity of the master problem.

The classic single-cut method performs more consistently than the multi-cut approach but remains significantly slower than the accelerated method, particularly at higher λ values. While single-cut avoids the explosion in cuts, it requires more iterations to converge, leading to moderate but non-negligible runtime. Overall, the proposed accelerated multi-cut strategy successfully combines the strengths of both classical methods: it avoids the instability and excessive computational burden of the multi-cut approach while achieving faster convergence than the single-cut method. This results in a substantial reduction in runtime, making the proposed method more suitable for large-scale and high-penalty scenarios.

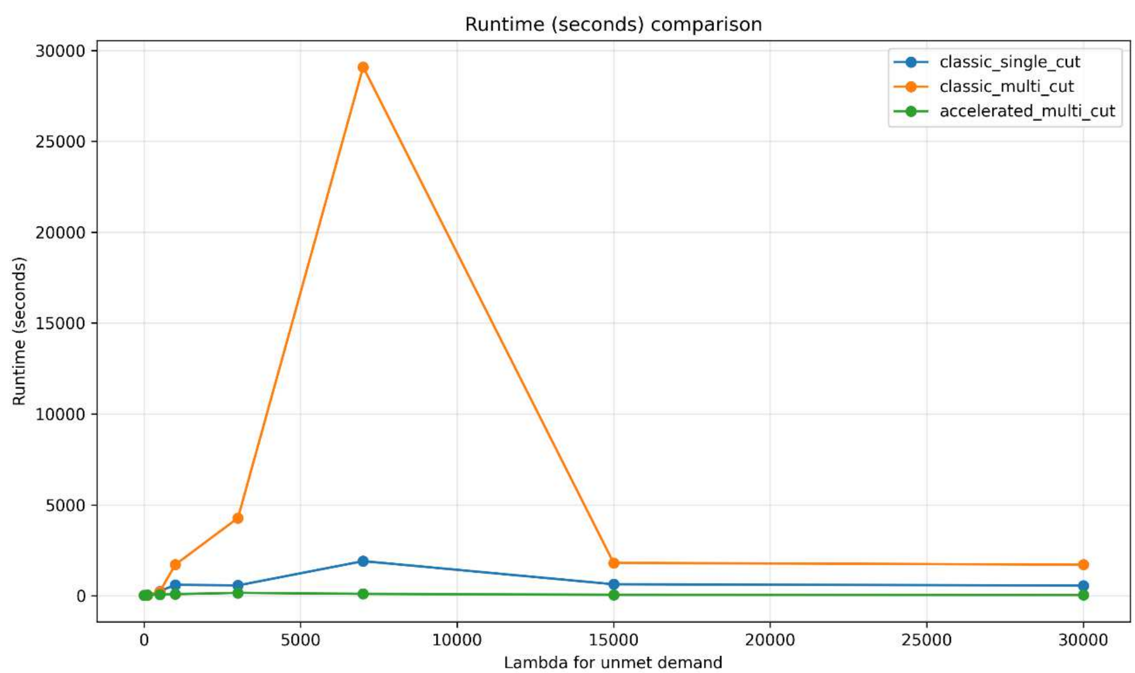

Figure 6 Runtime comparison for three approaches

Figures 7 (a-f) present the convergence behavior of the three Benders decomposition strategies—classic single-cut, classic multi-cut, and the proposed accelerated multi-cut—in terms of the relative optimality gap for different values of the unmet demand penalty parameter ($\lambda \in \{500, 1000, 3000, 7000, 15000, 30000\}$).

Across all tested scenarios, a consistent pattern emerges. First, the accelerated multi-cut method demonstrates the fastest convergence in every case. It rapidly reduces the optimality gap to near-zero within a very small number of iterations (typically fewer than 10 iterations), regardless of the value of λ. This behavior indicates that the proposed method is highly effective in generating informative cuts that significantly tighten the master problem early in the optimization process.

In contrast, the classic multi-cut approach shows improved convergence over the single-cut method, particularly in the early iterations, due to the addition of multiple cuts per iteration. However, this advantage diminishes as λ increases. Although the multi-cut method converges faster than single-cut, it still requires substantially more iterations than the accelerated approach, especially in high-penalty scenarios.

The classic single-cut method consistently exhibits the slowest convergence across all λ values. Its gap reduction is gradual and requires a significantly larger number of iterations (often exceeding 70–80 iterations for large λ values), reflecting the limited information provided by a single cut per iteration.

A key observation is that the performance gap between the accelerated method and classical approaches becomes more pronounced as λ increases. For higher penalty values (λ = 7000, 15000, 30000), the accelerated method maintains stable and rapid convergence, while both classical methods experience slower and more prolonged convergence processes.

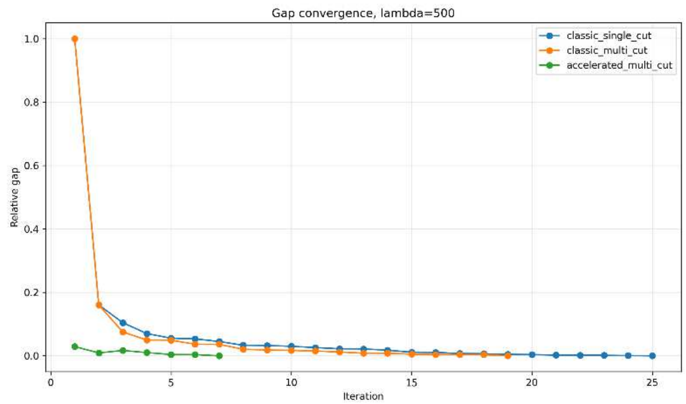

(a) Lambda=500

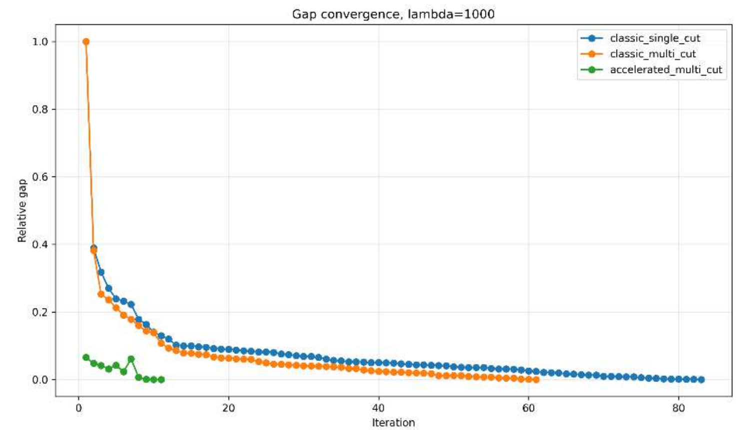

(b) Lambda=1000

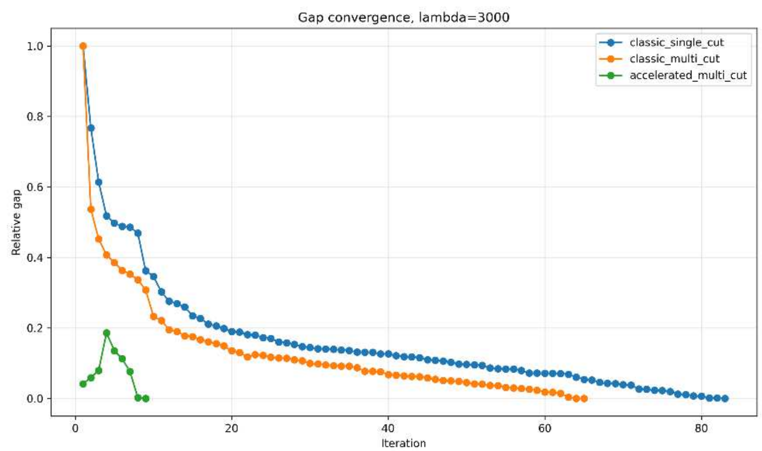

(c) Lambda=3000

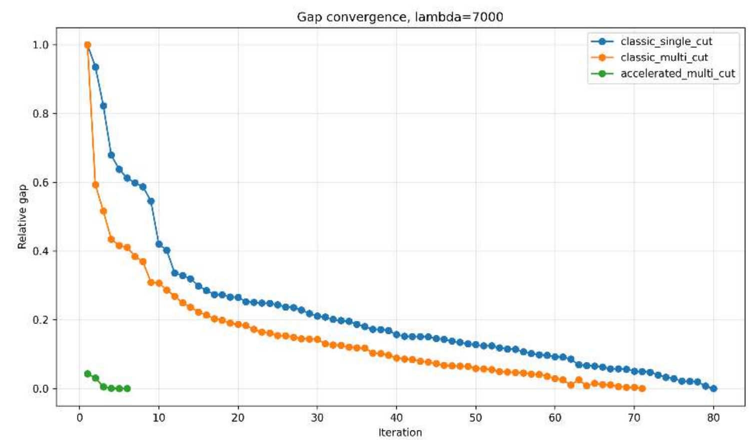

(d) Lambda=7000

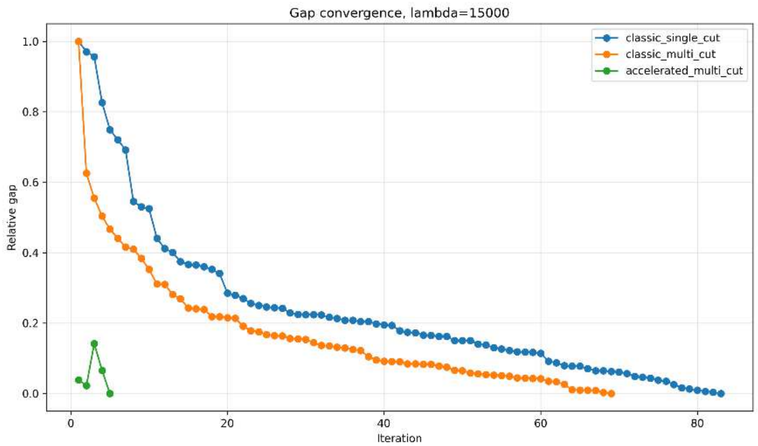

(e) Lambda=15000

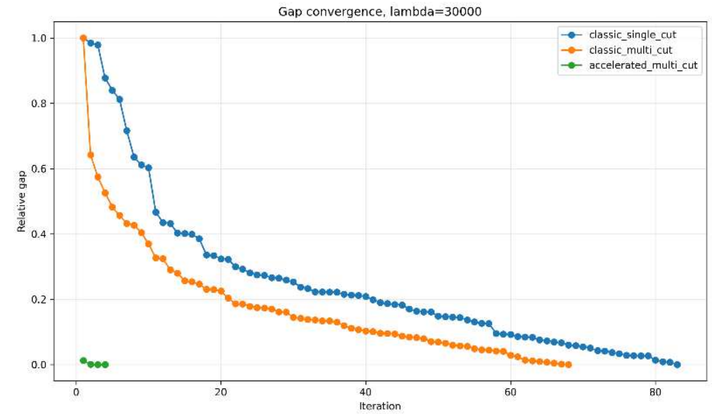

(f) Lambda=30000

Figure 7 Gap Convergence with different lambda values

Figure 8 (a-f) presents the LB/UB convergence with different lambda. The convergence behavior of the lower bound (LB) and upper bound (UB) across different values of the unmet demand penalty parameter (λ) demonstrates clear performance differences among the three solution strategies. Across all tested scenarios, the accelerated multi-cut approach consistently exhibits the fastest convergence, achieving near-optimal bounds within a very small number of iterations. In contrast, both the classic single-cut and classic multi-cut methods require substantially more iterations to close the optimality gap, particularly as λ increases.

For smaller values of λ (e.g., $\lambda = 500$ and 1000), all methods eventually converge to similar UB values; however, the accelerated multi-cut method rapidly stabilizes both LB and UB within the first few iterations, indicating strong early-stage bound tightening. The classic multi-cut method improves upon the single-cut approach by providing tighter bounds earlier, but still lags significantly behind the accelerated version. This gap becomes more pronounced as λ increases, where the problem becomes more constrained and computationally demanding.

At higher penalty levels ($\lambda = 3000$ to 30000), the superiority of the accelerated multi-cut method becomes even more evident. The LB trajectory of the accelerated approach quickly approaches the UB, while the classic methods exhibit slow, gradual improvements over many iterations. Notably, the classic single-cut method shows the slowest convergence, with a persistent gap between LB and UB even after a large number of iterations. The classic multi-cut method reduces this gap more effectively but still requires considerably more iterations compared to the accelerated variant.

From the results, it is confirmed that the proposed accelerated multi-cut framework significantly enhances convergence efficiency by improving bound quality early in the solution process and reducing the total number of iterations required to reach optimality.

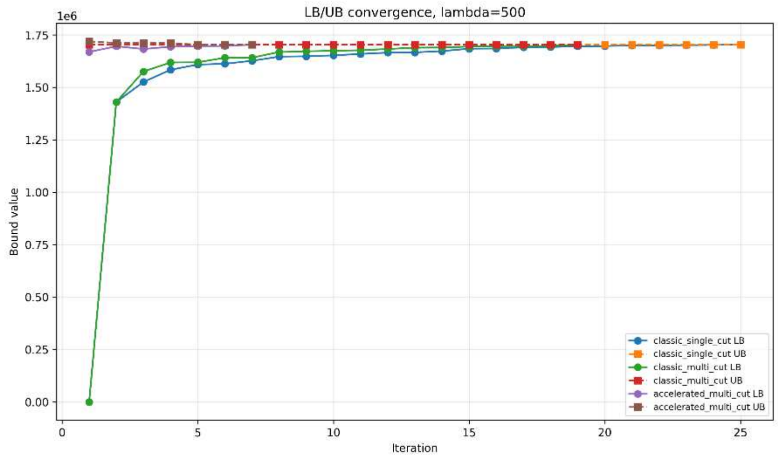

(a) Lambda=500

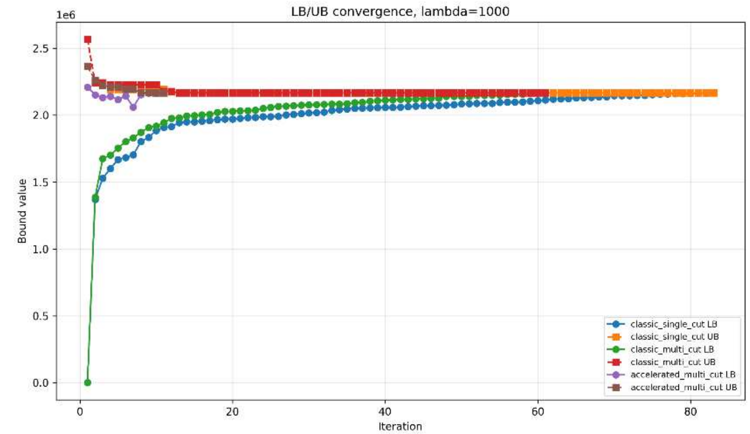

(b) Lambda=1000

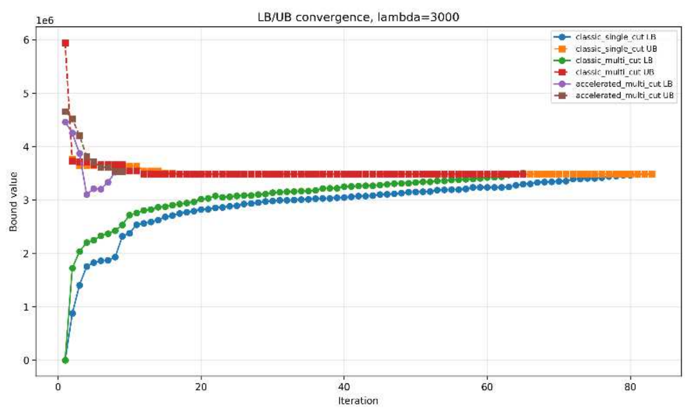

(c) Lambda=3000

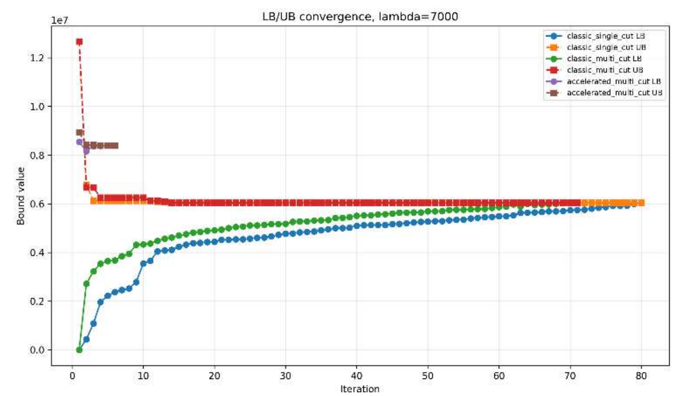

(d) Lambda=7000

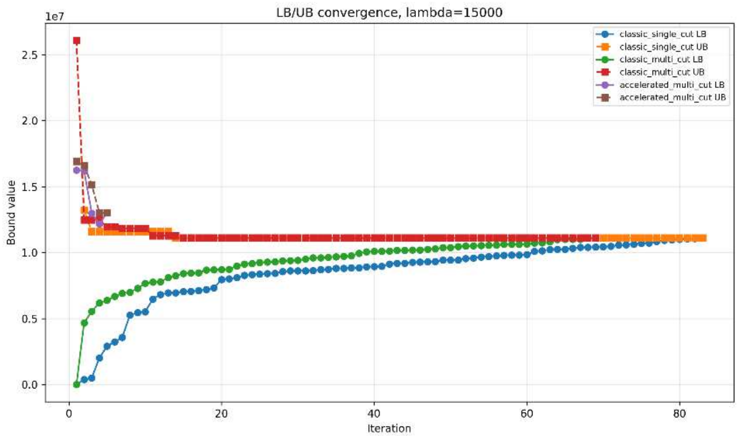

(e) Lambda=15000

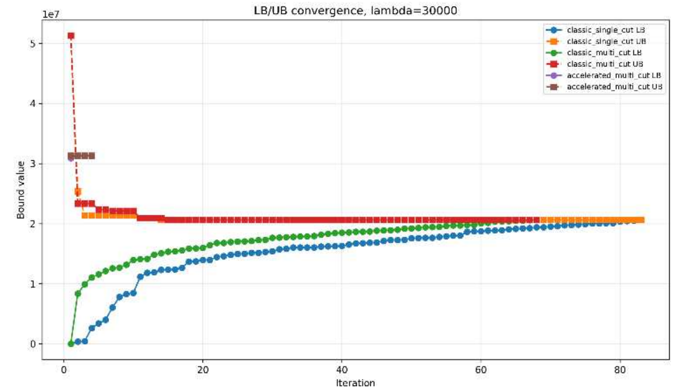

(f) Lambda=30000

Figure 8 LB/UB convergence with different lambda

6. Discussion

The obtained results demonstrate that resilient air transport network design under geopolitical disruption requires a careful balance between operational efficiency and service continuity. The proposed framework successfully captures this trade-off through the unmet-demand penalty parameter (λ), which directly controls the relative importance of resilience in the optimization process. As λ increases, the model progressively shifts from a cost-oriented strategy toward a resilience-oriented strategy, resulting in lower unmet demand but higher operational and reserve allocation costs. This behavior is consistent with real-world airline operations, where maintaining network continuity during severe disruptions often requires significant investments in rerouting flexibility, backup capacity, and reserve resources.

One important observation from the numerical experiments is that resilience improvements are not achieved through uniform reserve distribution across the network. Instead, the optimization model concentrates reserve capacity on a limited number of strategically important hubs, particularly DXB and IST, with secondary allocations to hubs such as AUH and CAI. This finding reflects the structural characteristics of hub-and-spoke aviation systems, where a small number of highly connected airports dominate traffic redistribution during disruptions. The results suggest that resilience-oriented investment strategies should prioritize critical connectivity hubs rather than attempting to reinforce all airports equally. Such targeted reinforcement can improve operational robustness while avoiding excessive infrastructure expenditure.

The hub utilization analysis further reveals that increasing resilience requirements leads to significantly higher utilization of major hubs under disruption scenarios. Several hubs approach near-capacity operating conditions for larger λ values, indicating that rerouted traffic becomes heavily concentrated on a small subset of airports. While this behavior improves service continuity and reduces unmet demand, it also introduces secondary operational risks, including congestion propagation, delay amplification, and reduced recovery flexibility. Therefore, the results highlight an important managerial implication: resilience enhancement may unintentionally create new vulnerability points if reserve allocation and traffic redistribution become overly concentrated.

Interestingly, some globally important airports, such as DOH, do not emerge as dominant reserve-allocation hubs despite maintaining moderate utilization levels in several scenarios. This suggests that resilience importance is not determined solely by airport size or international prominence, but rather by the interaction between network topology, rerouting flexibility, available capacity, and marginal recovery efficiency. In the analyzed network, DXB and IST provide stronger disruption absorption capability and more effective traffic redistribution potential under the considered disruption patterns. This observation highlights the importance of network-aware resilience planning rather than relying purely on traffic volume or geographical significance.

From an algorithmic perspective, the proposed accelerated multi-cut Benders decomposition demonstrates substantial computational advantages over the classical single-cut and classical multi-cut approaches. The results show that the accelerated strategy achieves significantly faster convergence, fewer iterations, and lower runtime across all tested λ values. The trust-region stabilization mechanism plays an important role in preventing oscillatory reserve-capacity decisions, while cut management reduces the excessive master-problem complexity commonly associated with classical multi-cut approaches. At the same time, the method preserves the scenario-specific information structure that is lost in aggregated single-cut formulations. Consequently, the proposed approach successfully combines the stability of single-cut decomposition with the stronger convergence characteristics of multi-cut methods.

The integration of VAE-generated scenarios also contributes important methodological value. Traditional stochastic optimization studies often rely on manually specified disruption scenarios, which may fail to capture the nonlinear and heterogeneous nature of geopolitical disruptions. By generating synthetic but realistic disruption patterns through latent-space learning, the proposed framework allows a richer representation of uncertainty while preserving computational tractability through scenario clustering. This learning-assisted scenario generation mechanism provides a flexible way to model evolving disruption conditions without relying exclusively on historical observations.

Despite these advantages, several limitations should be acknowledged. First, the disruption scenarios are partially based on semi-synthetic assumptions related to Middle East airspace disruptions rather than fully observed operational datasets. Although the VAE improves scenario diversity, the realism of the generated disruptions still depends on the quality and representativeness of the input data. Second, the model focuses primarily on strategic network reconfiguration and does not explicitly include detailed airline operational recovery decisions such as crew scheduling, aircraft rotation, or passenger reaccommodation. Third, the current formulation considers static disruption realizations within each scenario and does not model time-dependent recovery dynamics or sequential decision-making processes.

Future research can extend the framework in several directions. Dynamic multi-period formulations could better capture disruption propagation and recovery evolution over time. Additional resilience measures such as delay propagation, carbon emissions, passenger fairness, or operational risk exposure could also be integrated into the optimization model. Furthermore, graph neural networks, reinforcement learning, or adaptive online optimization methods may provide promising directions for real-time disruption management and large-scale aviation decision support under rapidly evolving geopolitical conditions.

7. Conclusion

This study proposed a resilient air transport network reconfiguration framework for managing large-scale geopolitical disruptions under uncertainty. The study introduced a bi-objective stochastic optimization model that simultaneously minimizes total operational cost and expected unmet demand while accounting for disruption-driven network reconfiguration. To represent uncertainty more realistically, a Variational Autoencoder (VAE) was employed to generate synthetic disruption scenarios capturing nonlinear relationships among capacity reductions, route closures, cost escalation, and demand variation. The generated scenarios were subsequently integrated into a stochastic optimization framework solved using an accelerated multi-cut Benders decomposition algorithm.

The computational results demonstrate that the proposed framework successfully captures the fundamental trade-off between operational efficiency and resilience. Increasing the unmet-demand penalty parameter leads to improved service continuity and reduced unmet demand, although at the cost of higher reserve allocation and operational expenditure. The results further show that resilience improvements are achieved primarily through strategic reinforcement of a small number of highly connected hubs rather than through uniform reserve allocation across the network. Major hubs such as DXB and IST emerge as dominant recovery and rerouting centers under severe disruptions, highlighting the structural importance of connectivity and network topology in resilient air transport systems.

From an algorithmic standpoint, the proposed accelerated multi-cut Benders decomposition substantially outperforms the classical single-cut and classical multi-cut methods in terms of runtime, convergence speed, and iteration efficiency. The combination of stabilization, cut management, and scenario-specific decomposition enables the framework to maintain computational tractability even for high-penalty and highly disrupted scenarios. The results confirm that learning-assisted decomposition methods can significantly enhance the scalability and practical applicability of stochastic network optimization models.

Overall, this study contributes to the growing literature on resilient transportation systems by bridging predictive disruption modeling and prescriptive optimization. The proposed framework provides a practical decision-support tool for airlines, airport operators, and transportation planners seeking to improve network adaptability under large-scale geopolitical uncertainty. Beyond aviation, the proposed methodology may also be extended to other critical infrastructure systems characterized by uncertainty, network interdependence, and disruption-driven reconfiguration requirements.

Conflict of Interest

No potential conflict of interest was reported by the author(s).

Funding

No funding was received for this study.

Data availability statement

The data that support the findings of this study are available from the corresponding author, Iman Rahimi, upon reasonable request.

References

Bai, R., Wallace, S. W., Li, J., & Chong, A. Y. L. (2014). Stochastic service network design with rerouting. *Transportation Research Part B: Methodological*, *60*, 50-65.

Carbonera, M., Ciavotta, M., & Messina, E. (2024). Variational autoencoders and generative adversarial networks for multivariate scenario generation. *Data Science for Transportation*, *6*(3), 23.

Chuong, T. D., Liu, C., & Yu, X. (2025). Decomposition for Large-Scale Optimization Problems: An Overview. *Artificial Intelligence Science and Engineering*, *1*(3), 157-174.

Crainic, T. G., Hewitt, M., Maggioni, F., & Rei, W. (2021). Partial benders decomposition: general methodology and application to stochastic network design. *Transportation Science*, *55*(2), 414-435.

Dong, W., Chen, X., & Yang, Q. (2022). Data-driven scenario generation of renewable energy production based on controllable generative adversarial networks with interpretability. *Applied Energy*, *308*, 118387.

Filho, A. A., Correia, I., & Melo, T. (2026). Bi-objective stochastic supply chain network design under supplier and facility disruption risks: AA Filho et al. *Computational and Applied Mathematics*, *45*(7), 309.

Gopalakrishnan, K., & Balakrishnan, H. (2021). Control and optimization of air traffic networks. *Annual Review of Control, Robotics, and Autonomous Systems*, *4*(1), 397-424.

Gui, G., Liu, F., Sun, J., Yang, J., Zhou, Z., & Zhao, D. (2019). Flight delay prediction based on aviation big data and machine learning. *IEEE Transactions on Vehicular Technology*, *69*(1), 140-150.

Hassan, L. K., Santos, B. F., & Vink, J. (2021). Airline disruption management: A literature review and practical challenges. *Computers & Operations Research*, *127*, 105137.

Kim, Y. J., Choi, S., Briceno, S., & Mavris, D. (2016, September). A deep learning approach to flight delay prediction. In *2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)* (pp. 1-6). IEEE.

Magnanti, T. L., & Wong, R. T. (1981). Accelerating Benders decomposition: Algorithmic enhancement and model selection criteria. *Operations research*, *29*(3), 464-484.

Müllerklein, D., & Fontaine, P. (2025). Resilient transportation network design with disruption uncertainty and lead times. *European Journal of Operational Research*, *322*(3), 827-840.

Paskov, A. S. (2025). *Large-Scale Optimization using Reinforcement Learning, Dynamic Programming, and Column Generation* (Doctoral dissertation, Massachusetts Institute of Technology).

Rahimi, I. (2025). Deep Learning Decision Support System for Open-Pit Mining Optimisation: GPU-Accelerated Planning Under Geological Uncertainty. *arXiv preprint arXiv:2511.18296*.

Rahimi, I. (2025). A decision support system for open-pit mining optimisation with dynamic uncertainty and GPU-based parallel repair approach. *Expert Systems with Applications*, 129544.

Rahmaniani, R., Crainic, T. G., Gendreau, M., & Rei, W. (2018). Accelerating the Benders decomposition method: Application to stochastic network design problems. *SIAM Journal on Optimization*, *28*(1), 875-903.

Liang, J., & Tang, W. (2019). Sequence generative adversarial networks for wind power scenario generation. *IEEE Journal on Selected Areas in Communications*, *38*(1), 110-118.

Sun, X., & Wandelt, S. (2021). Robustness of air transportation as complex networks: Systematic review of 15 years of research and outlook into the future. *Sustainability*, *13*(11), 6446.

Su, Y., Xie, K., Wang, H., Liang, Z., Chaovalitwongse, W. A., & Pardalos, P. M. (2021). Airline disruption management: A review of models and solution methods. *Engineering*, *7*(4), 435-447.