engineering proceedings

MDPI

Proceeding Paper

Topological Segmentation of Structured Reticular Multicommodity Flow in NP-Complete Problems: An A Priori Family Evacuation Metamodel for Meteorological Events †

Yasmany Fernández-Fernández [1,*], Sira M. Allende-Alonso [2], Ridelio Miranda-Pérez [3], Gemayqzel Bouza-Allende [2] and Kevin L. Ortega Chávez [1,*]

[1] Carrera de Computación, Universidad Politécnica Estatal del Carchi (UPEC), Tulcán 040102, Ecuador
[2] Facultad de Matemática y Computación (MATCOM), Universidad de La Habana, Havana 10400, Cuba; sira@matcom.uh.cu (S.M.A.-A.); gema@matcom.uh.cu (G.B.-A.)
[3] Facultad de Ciencias Económicas, Universidad de Cienfuegos, Cienfuegos 55100, Cuba; rmiranda@ucf.edu.cu
* Correspondence: yfernandezf@upec.edu.ec (Y.F.-F.); kevin.ortega@upec.edu.ec (K.L.O.C.)
† Presented at the Sixth International Conference on Information Systems and Software Technologies (ICI2ST 2026), Quito, Ecuador, 23–25 September 2026.

Abstract

Hurricane evacuation requires two decisions that the literature usually studies apart—where to send people and with which vehicles to move them—both fixed 24 to 48 h before landfall and both exposed to different kinds of uncertainty: demand is random, since it follows a storm track known only in probability, while travel times are epistemically uncertain, since they depend on road degradation that admits no reliable distribution. Existing models address one source or the other, treat evacuees as an aggregate headcount, and rarely carry the information produced by the scenario analysis into the uncertainty set used for protection. This paper proposes a hierarchical metamodel that closes those three gaps. A deterministic layer routes families of heterogeneous size as distinct integer commodities over a topologically segmented network with endogenous shelter activation; a second layer schedules a heterogeneous fleet over the extracted routes under vehicle–origin authorization, route consolidation and makespan minimization. A two-stage stochastic extension replaces nominal demand with an ensemble of storm scenarios produced by an adaptive generator, and static and adaptive robust extensions immunize the schedule against travel-time perturbations through uncertainty sets anchored to that same ensemble. The decision version is proven NP-hard, while the surrounding pipeline remains polynomial, and a hierarchical decomposition by origin preserves optimality while reducing the binary decision space by 93.5%. On a case study of Pinar del Río (Cuba) under hurricane IAN (2022), the deterministic core solves to proven optimality in under a second and the full robust cascade in about six minutes, returning diagnostics directly usable by civil-defense planners.

Keywords: evacuation planning; humanitarian logistics; multicommodity network flow; hierarchical decomposition; robust optimization; vehicle scheduling

Check for updates

Academic Editor: Thomás Borja

Published: 18 September 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Emergency evacuation planning under meteorological threats combines integer multicommodity network flow optimization, endogenous shelter activation and vehicle fleet scheduling under severe pre-impact time constraints. Hurricanes and other forecastable extreme events impose a particular operational discipline: shelter activation, route allocation and fleet deployment must be decided 24 to 48 h beforelandfall, when the track,

the intensity and the realized demand are still uncertain [1,2]. The Caribbean and the Gulf of Mexico, exposed every year to tropical cyclones with significant socio-economic impact, make this class of problems both academically rich and operationally critical [3,4].

The Cuban civil-defense system, which motivates this work, illustrates the structure of the decision. The provincial Defense Council identifies the vulnerable families and their embarkation points a priori from the projected track, i.e., the rainfall and the coastal flooding maps; it knows the evacuation centers, along with their capacities and activation costs; and it operates a state fleet whose composition is known in the preparation phase but whose operating times on a degraded road network are not. The problem is therefore, by construction, an optimization problem under two heterogeneous sources of uncertainty: demand is random because the track is known only probabilistically, and travel times are epistemically uncertain because road degradation and evacuee behavior admit no credible probabilistic characterization.

Three shortcomings of the current literature follow. First, the two dominant paradigms are applied in isolation: two-stage stochastic programming [5,6] captures demand variability through scenarios with explicit probabilities, robust optimization [7–9] immunizes decisions against bounded perturbations, and few works combine them on the same problem—let alone propagate information from the scenario ensemble into the geometry of the uncertainty set. Second, network routing and fleet scheduling are, in the large majority of contributions, solved separately, so the transformation from an optimal arc-flow pattern into an executable set of vehicle trips has no formal treatment. Third, evacuees are modeled as an aggregate headcount, which is convenient mathematically and wrong operationally: a bus cannot take three members of a family of five and leave the rest behind. Section 2 documents these gaps against the recent literature and against a systematic bibliometric analysis of 6046 records indexed in Scopus for 2022–2026.

The framework has three vertically coupled layers. A deterministic flow layer (P_I^D) operates over a topologically segmented reticular network: families are grouped by size and treated as distinct integer commodities, transit and shelter nodes are activated endogenously, and arc and node capacities are honored simultaneously [10,11]. An origin-authorized scheduling layer (P_{II}^D) takes the routes extracted from P_I^D and schedules a heterogeneous fleet under consolidation, sequencing and makespan minimization. Above them, the stochastic extension (P_I^S) replaces nominal demand with a scenario family produced by an adaptive hurricane scenario generator (MGEA), and the robust extensions (P_{II}^R, P_{II}^{ARO}) protect the schedule against travel-time perturbations through uncertainty sets anchored to that ensemble.

What distinguishes the metamodel from existing evacuation models can be stated in five claims, substantiated in Section Research Gaps and Positioning and summarized in Table 1: (N1) *family-level granularity as a modeling primitive*—family size is a commodity attribute, not an aggregation weight, so no family is split across vehicles or shelters, a property aggregate formulations [3,4,12] cannot express; (N2) a *formal flow-to-schedule transformation*, whose output partition, by origin, is not incidental but the structural artifact that makes the scheduling layer decomposable; (N3) a *decomposition with provable optimality preservation* (Theorem 1) rather than a heuristic partition; (N4) *scenario-anchored uncertainty sets*, whose amplitude is inherited from the scenario posterior instead of the customary uniform inflation factor; and (N5) a *complexity characterization of the coupled* flow-plus-scheduling decision problem, not of an isolated component, with the boundary between hard cores and polynomial pipeline made explicit.

Section 2 reviews the literature and positions the contribution; Sections 3–5 develop the flow layer, the vehicle layer and its decomposition, and the stochastic and robust extensions; Section 6 covers complexity, the rationale for exact rather than metaheuristic solution and

https://doi.org/10.3390/engproc2026153007

the feasibility of the sequential pipeline; Section 7 reports the case study, as well as its verification and validation, and comparison with the state of the art; and Sections 8 and 9 present the implementation pathway and the limitations of this study.

2. Literature Review and Positioning

The review is informed by a systematic bibliometric analysis of Scopus for 2022–2026 under the protocol of Kitchenham and Charters [13], with seven progressively restrictive queries and a corpus of 6046 records. Three of its measurements recur below: stochastic programming applied to evacuation returns 3237 records against 73 for static robust optimization on the same domain; a qualitative screening finds fewer than fifteen publications applying both paradigms to the same problem; and the intersection between robust optimization, evacuation and indeterminacy-based calibration of uncertainty sets is empty.

Network-flow models and shelter location. Evacuation entered operations research as a network flow problem, and the surveys of Hamacher and Tjandra [1] and Murray-Tuite and Wolshon [2] remain the standard entry points. Location and routing were later integrated: Bayram and Yaman [4] solved shelter location with route assignment under uncertainty by Benders decomposition, Esposito Amideo et al. [14] reviewed the critical issues of that coupling, and Yin et al. [3] proposed a two-stage recoverable robust model. All share two features relevant here: demand is an aggregate attached to a node, and the output is a flow pattern, not a vehicle assignment.

Transit-based evacuation and vehicle scheduling. A parallel and less crowded line studies populations without private vehicles, which is the dominant case in the Caribbean. Swamy et al. [15] planned hurricane evacuation with public transportation; Zhao et al. [16] formulated round-trip bus evacuation with joint scheduling and routing; Goerigk et al. [17] solved a comprehensive planning model using a genetic algorithm; Dulebenets et al. [18] and Khayamim et al. [19] assigned individuals, including vulnerable groups, to shelters through designated routes and time periods—the latter under pandemic conditions with a bi-objective model. This literature schedules vehicles competently but takes routing as given or solves both decisions monolithically.

Uncertainty: stochastic and robust paradigms. Two-stage and multi-stage stochastic programming dominates humanitarian logistics [5,6]: Rawls and Turnquist [20] set the template for pre-positioning; Vahdani et al. [21] combined location, routing and inventory; Karabuk and Manzour [22] addressed track uncertainty; Siddig and Song [23] studied multi-stage hurricane evacuation; and Heng et al. [24] integrated facility location, inventory, evacuation and vehicle logistics with scenarios derived from flood hazard maps. The scenario ensemble is the critical input of all of them and also their weakest link, since probabilities are frequently uniform or expert-assigned. Robust optimization replaces distributions with uncertainty sets [7–9], and affine decision rules [25] recover part of the conservatism that worst-case immunization [26] gives away; in evacuation, the works of Eriskin and Karatas [12] and Yin et al. [3] are representative. In every reviewed robust evacuation model, the amplitude of the set is a uniform inflation factor—25% is the recurring choice—justified only by convention.

Solution strategies. Because these models are NP-hard, the field has largely turned to approximate methods, such as genetic algorithms [17], with heuristics adopted after exact methods proved too slow [18], as well as decomposition-based ε-constraint procedures [19]; exact decomposition has also succeeded, with Benders in [4] being the clearest case. Metaheuristics scale but return no optimality certificate, which is uncomfortable when the approximated object is the number of people left behind. Section 6.2 explains why we exploit structure to keep exact solvers viable instead.

https://doi.org/10.3390/engproc2026153007

ResearchGaps and Positioning

Six mutually reinforcing gaps emerge: (G1) fragmented treatment of the two uncertainty paradigms, with fewer than 2% of the stochastic corpus also applying robust optimization to the same problem; (G2) ad hoc uncertainty sets, with no mechanism anchoring them to the available scenario information; (G3) absent hierarchical integration of flow and scheduling, with the flow-to-route transformation having no formal treatment; (G4) aggregate demand granularity, leaving family indivisibility outside the model; (G5) uncharacterized complexity of the integrated system, with NP-hardness being established only for isolated components; and (G6) scalability barriers, with monolithic formulations degrading beyond fifty nodes and decompositions that do not preserve the semantics of the layers.

Table 1 places this work against representative contributions along those dimensions. The pattern, rather than any single missing feature, is what the metamodel addresses: each existing work covers two or three dimensions well, and the columns for family granularity, formal flow-to-schedule coupling and complexity analysis of the integrated problem are essentially empty.

Table 1. Positioning against representative contributions. SP = stochastic programming; RO = robust optimization; ADR = affine decision rule; Sch. = fleet scheduling; Fam. = family-level granularity; Dec. = exact decomposition; Cpx. = complexity of the integrated problem. ● substantial; ◦ partial; — not addressed.

Study	Problem Addressed	SP	RO	ADR	Sch.	Fam.	Dec.	Cpx.
Bayram & Yaman [4]	Shelter location + routing	●	—	—	—	—	●	—
Goerigk et al. [17]	Comprehensive evacuation	◦	◦	—	●	—	—	◦
Eriskin & Karatas [12]	Hurricane evacuation	—	●	—	◦	—	—	—
Yin et al. [3]	Location allocation + evacuation	●	●	◦	—	—	◦	—
Heng et al. [24]	Location, inventory, and evacuation	●	—	—	◦	—	—	—
Khayamim et al. [19]	Evacuation under pandemic conditions	—	—	—	◦	◦	●	—
Dulebenets et al. [18]	Evacuation of vulnerable groups	—	—	—	◦	◦	—	—
This work	Flow + fleet, hierarchical	●	●	●	●	●	●	●

3. Problem Setting, Notation and the Deterministic Flow Layer

Table 2 collects every symbol used in the paper so that each formulation below can be read against a single reference; each block of equations is then explained constraint by constraint.

Table 2. Notation: persons and families are integers; costs in U.S. dollars; times in hours.

Symbol	Meaning	Symbol	Meaning
$\mathcal{G} = (\mathcal{N}, \mathcal{A}, \mathcal{D})$	network, arcs, distance matrix	Cap_v	capacity of vehicle v (persons)
$\mathcal{N}_A, \mathcal{N}_R, \mathcal{N}_F$	origins, transit nodes, shelters	c_v^f, c_v^v	fixed ($) and per-km ($/km) cost of v
$\mathcal{F}_1, \ldots, \mathcal{F}_T$	topological segments of $\mathcal{N}_F$	ρ_v	return-to-origin indicator of v
$H = \{1, \ldots, h_{\max}\}$	admissible family sizes	$\mathrm{VehOrig}_{v,o}$	authorization of v at origin o
$\mathcal{I}, \mathcal{I}_o$	family identifiers; those with origin o	d_r, f_r, h_r, o_r	length, demand, size class, origin of r
Ω, Ω_o	extracted routes; those with origin o	$v_{\mathrm{med}}, \tau_{\mathrm{serv}}$	fleet speed; per-person service time
$\mathcal{U}, \mathcal{U}_o$	fleet; vehicles assigned to o	$T_{\max}, w_{\mathrm{ms}}$	pre-impact horizon; makespan weight
$\mathcal{K} = \{1, \ldots, K_{\max}\}$	trip indices per vehicle	ε, M_μ	per-trip and coverage-slack penalties
$\mathcal{S}$	meteorological scenarios	$f^s(h,i), p_s$	scenario demand and probability
d_{ij}, c	arc length (km); cost ($/person-km)	θ^+, θ^-	deficit/over-evacuation penalty
$f^0(h,i), P_0$	nominal families at i; total demand	$\bar{\tau}_r, \rho^\tau$	nominal route duration; deviation

Table 2. *Cont.*

Symbol	Meaning	Symbol	Meaning
$q_{\mathrm{id}}, h_{\mathrm{id}}, o_{\mathrm{id}}$	families, size and origin of identifier	$\sigma(r), \Gamma$	anchored deviation; robustness budget
u_{ij}, α_i	arc capacity; embarkation capacity	$x_{\mathrm{id},i,j} \in \mathbb{Z}_+$	families of id on arc (i,j)
β_j, π_j	transit throughput; shelter capacity	$y_j \in \{0,1\}$	activation of node j
K_j, λ_k	activation cost; capacity reserve factor	$W_{v,r,k} \in \{0,1\}$	v serves r as its k-th trip
$\varphi_{\mathrm{id},i}$	net supply of id at node i	$F_{v,r,k} \in \mathbb{Z}_+$	families carried on that trip
$a_{\mathrm{id},i}$	families of id sheltered at i	$U_v, Q_{v,k} \in \{0,1\}$	vehicle activation; trip activity
$\mathcal{U}_{\mathrm{box}}, \mathcal{U}_\Gamma^\delta$	box and budgeted uncertainty sets	$T^{\mathrm{ini}}_{v,k}, T^{\mathrm{fin}}_{v,k}$	start and end time of trip k of v
$\mu_r \in \mathbb{Z}_+$	uncovered demand on route r	T^{ms}, $\delta^{\pm}_{s,i}$	makespan; second-stage deficit/excess

3.1. Reticular Network, Family Commodities and Segmentation

Let $\mathcal{G} = (\mathcal{N}, \mathcal{A}, \mathcal{D})$ be a directed graph whose node set is partitioned topologically, i.e.,

$$\mathcal{N} = \mathcal{N}_A \cup \mathcal{N}_R \cup \mathcal{N}_F, \qquad \mathcal{N}_A \cap \mathcal{N}_R = \mathcal{N}_A \cap \mathcal{N}_F = \mathcal{N}_R \cap \mathcal{N}_F = \emptyset, \tag{1}$$

where $\mathcal{N}_A$ denotes populated origins (municipalities at risk), $\mathcal{N}_R$ denotes transit nodes and $\mathcal{N}_F$ denotes shelter centers; $\mathcal{D} = (d_{ij})$ is derived from the actual road network [11]. Families are grouped by size into $H = \{1, \ldots, h_{\max}\}$, with $h_{\max} \leq 8$ in the case study, and $f^0(h,i)$ denotes the nominal number of families of size h at origin i so that the total demand is $P_0 = \sum_{h,i} h f^0(h,i)$ persons. Family size is treated as a *commodity attribute* rather than aggregated into headcount: this is essential operationally, since a vehicle cannot split a family across trips, and mathematically, since it gives family-level traceability throughout the metamodel.

Following [11], the demand is encoded into a set $(\mathcal{I})$ of *family identifiers*, each representing q_{id} families of size h_{id} originating at o_{id}, with $P_0 = \sum_{\mathrm{id}} h_{\mathrm{id}} q_{\mathrm{id}}$. The identifier set admits the partition of $\mathcal{I} = \bigsqcup_{o \in \mathcal{N}_A} \mathcal{I}_o$, the bookkeeping device that propagates through the pipeline and enables the decomposition of the vehicle layer (Theorem 1).

Definition 1 (Topological segmentation)**.** *A topological segmentation is a partition* $(\mathcal{N}_F = \mathcal{F}_1 \sqcup \cdots \sqcup \mathcal{F}_T)$ *of the shelter set into T disjoint, non-empty segments, each associated with (i) a subset of origins whose flow is structurally directed to it and (ii) a subset of transit nodes that primarily feed it.*

Segmentation operates between the individual shelter and the whole province, at the scale at which civil-defense capacity policies are set: exposure is approximately homogeneous within a segment, so scenario impact factors apply uniformly, road connectivity preserves intra-segment dominance and per-segment capacity rules preserve territorial reserve margins. The case study partitions Pinar del Río into western, central and eastern modules (Figure 1).

3.2. Deterministic Flow Layer $(\mathrm{P}_{\mathrm{I}}^{\mathrm{D}})$

With $x_{\mathrm{id},i,j} \in \mathbb{Z}_+$ representing the families of identifier id routed on arc (i,j) and $y_j \in \{0,1\}$ corresponding to the activation of node j, the deterministic flow problem is

$$\min_{x,y} \sum_{\mathrm{id}\in\mathcal{I}} \sum_{(i,j)\in\mathcal{A}} c\, d_{ij}\, h_{\mathrm{id}}\, x_{\mathrm{id},i,j} + \sum_{j\in\mathcal{N}_R\cup\mathcal{N}_F} K_j\, y_j, \tag{2a}$$

$$\text{s.t.} \quad \sum_{j:(i,j)\in\mathcal{A}} x_{\mathrm{id},i,j} - \sum_{j:(j,i)\in\mathcal{A}} x_{\mathrm{id},j,i} = \varphi_{\mathrm{id},i}, \quad \forall\, \mathrm{id}\in\mathcal{I},\ \forall\, i\in\mathcal{N}, \tag{2b}$$

$$\sum_{\mathrm{id}\in\mathcal{I}} h_{\mathrm{id}}\, x_{\mathrm{id},i,j} \le u_{ij}, \quad \forall\,(i,j)\in\mathcal{A}; \qquad \sum_{\mathrm{id}:\, o_{\mathrm{id}}=i} h_{\mathrm{id}} \sum_{j:(i,j)\in\mathcal{A}} x_{\mathrm{id},i,j} \le \alpha_i, \quad \forall\, i\in\mathcal{N}_A, \tag{2c}$$

$$\sum_{\mathrm{id}\in\mathcal{I}} \sum_{i:(i,j)\in\mathcal{A}} h_{\mathrm{id}}\, x_{\mathrm{id},i,j} \le \pi_j\, y_j \ \ \forall\, j\in\mathcal{N}_F; \qquad \sum_{\mathrm{id}\in\mathcal{I}} \sum_{i:(i,j)\in\mathcal{A}} h_{\mathrm{id}}\, x_{\mathrm{id},i,j} \le \beta_j\, y_j \ \ \forall\, j\in\mathcal{N}_R, \tag{2d}$$

$$x_{\mathrm{id},i,j}\in\mathbb{Z}_+, \qquad y_j\in\{0,1\}, \tag{2e}$$

where the type-dependent net supply is $\varphi_{\mathrm{id},i} = +q_{\mathrm{id}}$ if $i = o_{\mathrm{id}} \in \mathcal{N}_A$, 0 if $i \in \mathcal{N}_R$ and $-a_{\mathrm{id},i}$ if $i \in \mathcal{N}_F$, with $\sum_{i\in\mathcal{N}_F} a_{\mathrm{id},i} = q_{\mathrm{id}}$.

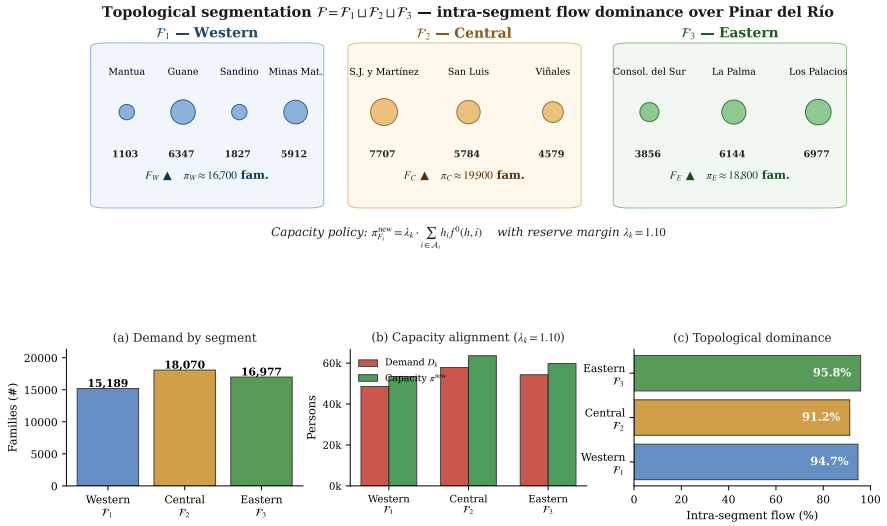


Figure 1. Segmentation, i.e., $\mathcal{F} = \mathcal{F}_1 \sqcup \mathcal{F}_2 \sqcup \mathcal{F}_3$: family demand by origin grouped by segment. (**a**) Demand by segment; (**b**) capacity alignment under the reserve policy of $\pi_j^{\mathrm{new}} = \lambda_k \pi_j^{\mathrm{base}}, \lambda_k = 1.10$; (**c**) measured intra-segment flow dominance at the $\mathrm{P}_{\mathrm{I}}^{\mathrm{D}}$ optimum.

The model reads as follows. Objective (2a) adds the transport cost—each family of identifier id on arc (i, j) contributes its h_{id} members over d_{ij} kilometers at the unit rate (c), so the term is proportional to person kilometers—and the fixed cost of opening transit nodes and shelters, which is what makes activation endogenous: a shelter opens only when the transport savings exceed the cost of running it. Constraint (2b) is flow conservation *per identifier*: at its origin, the net outflow equals the q_{id} represented families; and at transit nodes, inflow and outflow balance so that families are neither created nor absorbed midway; and at shelters, the net inflow equals the families finally sheltered there, with the condition of $\sum_i a_{\mathrm{id},i} = q_{\mathrm{id}}$ ensuring that nobody is left unassigned. The first inequality in (2c) limits the *persons* traversing an arc during the planning window, u_{ij} being practical throughput under evacuation conditions rather than free-flow capacity; the second bounds, by α_i, the persons that can be gathered and embarked at an assembly point. Both inequalities in (2d) carry the activation variable on the right-hand side and therefore do two jobs at once: they cap the persons received by node j, and they force $y_j = 1$—paying K_j—as soon as any flow enters, since a positive left-hand side is infeasible when $y_j = 0$. Thus, the transit nodes are *not* unlimited waypoints: β_j bounds how many people can pass through, be temporarily sheltered or be transferred there within the pre-impact window, and ignoring it would produce plans that funnel an entire province through one convenient set of crossroads. Finally, (2e) declares integrality, which is substantive rather than a convenience: rounding a fractional family has no operational meaning. The result is an integer multicommodity

flow problem with binary node activations and three simultaneous capacity levels—arc, origin and node.

3.3. Polynomial-Time Route Extraction

$\mathrm{P}_{\mathrm{I}}^{\mathrm{D}}$ yields aggregate arc flows. Algorithm 1, a depth-first search on the support graph of the positive flows, extracts explicit family routes in $O(|\mathcal{J}|(|\mathcal{N}| + |\mathcal{A}|))$ time; the partition of $\Omega = \bigsqcup_o \Omega_o$ it returns is the structural artifact enabling the decomposition of Section 4. The procedure terminates because each inner iteration removes at least one unit of flow, and it is exact in the sense that the extracted routes reconstruct the optimal arc flows: for every arc, the sum of f_r over the routes containing it equals x^* on that arc, an identity checked numerically for every instance (Section 7.3).

Algorithm 1 Family route extraction (DFS over the $\mathrm{P}_{\mathrm{I}}^{\mathrm{D}}$ support)

Require: Optimal x^* of $\mathrm{P}_{\mathrm{I}}^{\mathrm{D}}$; identifiers $\mathcal{J}$; partition $\mathcal{N} = \mathcal{N}_A \sqcup \mathcal{N}_R \sqcup \mathcal{N}_F$.
Ensure: Routes Ω, demands (f_r), assignment σ_r.
1: $\Omega \leftarrow \emptyset$, $\sigma_r \leftarrow \emptyset$.
2: **for** each id $\in \mathcal{J}$ with $q_{\mathrm{id}} > 0$ **do**
3: Build $\mathcal{G}_{\mathrm{id}}^{+} = \{(i,j) \in \mathcal{A} : x^*_{\mathrm{id},i,j} > 0\}$ and run DFS from o_{id}.
4: **for** each elementary path $r = (i_1, \dots, i_\ell)$, $i_1 = o_{\mathrm{id}}$, $i_\ell \in \mathcal{N}_F$ **do**
5: $f_r \leftarrow$ min capacity along r; $\Omega \leftarrow \Omega \cup \{r\}$; record h_r, o_r; decrement x^* by f_r along r.
6: **end for**
7: **end for**
8: **return** $(\Omega = \bigsqcup_{o \in \mathcal{N}_A} \Omega_o, \{f_r\}, \sigma_r)$.

4. Origin-Authorized Vehicle Scheduling and Hierarchical Decomposition

4.1. Vehicle Layer Formulation

The vehicle layer receives the routes (Ω), their demands $(\{f_r\})$ and labels $(\{h_r, o_r\})$. The objective combines fixed activation costs, variable transport costs, makespan weighted by w_{ms}, a small per-trip penalty (ε) and a heavy slack penalty (M_μ):

$$\min \sum_{v \in \mathcal{V}} c_v^f U_v + \sum_{v,r,k} c_v^v d_r \left(2\rho_v + (1-\rho_v)\right) W_{v,r,k} + w_{\mathrm{ms}} T^{\mathrm{ms}} + \varepsilon \sum_{v,k} Q_{v,k} + M_\mu \sum_{r \in \Omega} h_r \mu_r, \quad (3)$$

subject to

$$\sum_{v,k} F_{v,r,k} + \mu_r \geq f_r \quad \forall\, r; \qquad F_{v,r,k} \leq \mathrm{Cap}_v\, W_{v,r,k} \quad \forall\, v,r,k, \quad (4a)$$

$$W_{v,r,k} \leq \mathrm{VehOrig}_{v,o_r} \quad \forall\, v,r,k; \qquad \textstyle\sum_{r \in \Omega} W_{v,r,k} \leq Q_{v,k} \leq U_v \quad \forall\, v,k, \quad (4b)$$

$$T^{\mathrm{fin}}_{v,k} \geq T^{\mathrm{ini}}_{v,k} + \textstyle\sum_{r \in \Omega} \tau_{v,r} W_{v,r,k} \quad \forall\, v,k; \qquad T^{\mathrm{ini}}_{v,k+1} \geq T^{\mathrm{fin}}_{v,k} \quad \forall\, v,\ k < K_{\max}, \quad (4c)$$

$$T^{\mathrm{fin}}_{v,k} \leq T_{\max} \quad \forall\, v,k; \qquad T^{\mathrm{ms}} \geq T^{\mathrm{fin}}_{v,k} \quad \forall\, v,k, \quad (4d)$$

where $\tau_{v,r} = \rho_v \tau_r^{\mathrm{ret}} + (1-\rho_v)\tau_r^{\mathrm{nr}}$, with $\tau_r^{\mathrm{ret}} = 2d_r/v_{\mathrm{med}} + \mathrm{Cap}_v \tau_{\mathrm{serv}}$ and $\tau_r^{\mathrm{nr}} = d_r/v_{\mathrm{med}} + \mathrm{Cap}_v \tau_{\mathrm{serv}}$.

In (3), the first term pays for every mobilized vehicle, discouraging activation of the whole fleet when a subset suffices; the second charges kilometers, with charges doubled for vehicles that must return $(\rho_v = 1)$; the third prices the makespan, turning a cost-minimizing schedule into a time-aware one; the fourth breaks ties between equivalent schedules and removes spurious empty trips; and the last penalizes, weighted by family size, any demand left uncovered. The intended hierarchy is $M_\mu \gg w_{\mathrm{ms}} \gg c_v^f, c_v^v \gg \varepsilon$, so the model never trades uncovered people for money. Constraint (4a) requires the families assigned to route r, plus the slack (μ_r), to reach its demand: the slack is not a luxury, since an infeasible model would tell the planner nothing, whereas a positive μ_r says exactly where and by how much the

fleet falls short; the second inequality allows a trip to carry families only if it is performed, and never above vehicle capacity, which is where family indivisibility becomes physical. In (4b), vehicle v may serve r only if authorized at that route's origin—the VehOrig matrix encodes the real pre-impact allocation of transport to municipalities, and it is also what confines each vehicle to one origin block, while the $W \le Q \le U$ chain allows one route per trip index and propagates the fixed cost consistently. Constraint (4c) builds the per-vehicle timeline, and (4d) encodes the meteorological deadline—no trip finishing after $T_{\max}$—while bounding T^{ms} below by every trip end, so minimizing it minimizes the completion time of the evacuation.

4.2. Route Consolidation and Decomposition by Origin

Definition 2 (Route containment)**.** *For $r, r' \in \Omega$ with node sequences $\sigma_r, \sigma_{r'}$, r is* contained *in r' ($r \subseteq r'$) if $o_r = o_{r'}$ and σ_r is a prefix of $\sigma_{r'}$;* $\mathrm{Cont}(r) = \{r' \in \Omega : r \subseteq r'\}$.

A vehicle serving a longer route physically traverses every node of the shorter one and therefore satisfies part of its demand without an additional trip. This is encoded by replacing the capacity inequality of (4a) with

$$F_{v,r,k} \le \mathrm{Cap}_v\Big(W_{v,r,k} + \sum_{r' \in \mathrm{Cont}(r)\setminus\{r\}} W_{v,r',k}\Big), \qquad \forall\, v, r, k, \tag{5}$$

which avoids artificially duplicated trips for nested flows on a shared backbone. Containment is defined only between routes with the same origin, so (5) stays inside a single origin block.

The monolithic problem, with $|\mathcal{V}||\Omega|K$ binaries, contains bin-packing as a special case and is NP-hard [27,28].

Theorem 1 (Hierarchical Decomposition by Origin)**.** *Let $\Omega_o = \{r : o_r = o\}$, let (Y^*, σ^*) solve the master problem that assigns vehicles to origins, and let $\mathcal{V}_o = \{v : Y^*_{v,o} = 1\}$ and $K_o \le K_{\max}$. Then, (i) the optimum of $\mathrm{P}^{\mathrm{D}}_{\mathrm{II}}$ is reconstructed exactly as the union of the optima of the per-origin subproblems defined over $\mathcal{V}_o \times \Omega_o \times \{1, \dots, K_o\}$, and (ii) $\sum_{o \in \mathcal{N}_A} |\mathcal{V}_o||\Omega_o|K_o < |\mathcal{V}||\Omega|K_{\max}$, strictly whenever $K_o < K_{\max}$ for some origin or $\Omega_o \subsetneq \Omega$.*

Proof (sketch). Once Y^* is fixed, objective (3) is separable along $\Omega = \bigsqcup_o \Omega_o$ because consolidation involves only routes with a common origin (Definition 2) and authorization (4b) prevents cross-origin scheduling. Activation and continuity decouple along $\mathcal{V}_o$, and (4d) aggregates through $T^{\mathrm{ms}} = \max_o T^{\mathrm{ms}}_o$; the remaining constraints decompose blockwise. For (ii), $\sum_o |\mathcal{V}_o| \le |\mathcal{V}|$ and $\sum_o |\Omega_o| = |\Omega|$ give the bound, with strictness following from $K_o \le K_{\max}$ and $\Omega_o \subsetneq \Omega$. □

Implementation has two phases: a master mixed-integer program of size $|\mathcal{V}||\mathcal{N}_A|$, assigning each vehicle to, at most, one origin, and the restricted scheduling problems; across the case study, this yields a mean reduction of 93.5% of active binaries. Given the master assignment, the reconstruction is exact, so the pair (master and subproblems) is an exact reformulation under the origin-authorization structure, not an approximation whose gap would need estimating.

5. Meteorological Scenarios and Stochastic/Robust Extensions

Adaptive scenario generation. The stochastic layer receives an ensemble produced by an adaptive Meteorological Generator of Evacuation Activations (MGEA), where each $s \in \mathcal{S}$ encodes a plausible trajectory and yields demands ($f^s(h,i)$) and probabilities ($p_s > 0$) summing to one. It combines [11,29]: a historical record of tropical-cyclone tracks over

the Cuban archipelago from the HURDAT2 reanalysis of the National Hurricane Center; a polynomial feature map ($\Phi : \Xi \to \mathbb{R}^9$) with four intensity features, four geometry features and one kinematic feature; a Gaussian-mixture clustering with the number of scenarios selected by the Bayesian information criterion and bounded by $|\mathcal{S}| \le k^* = 6$; and a horizon-conditional reweighting that sharpens pessimistic clusters as lead time decreases, i.e.,

$$p_s^{\text{final}}(\tau) = \frac{p_s^{\text{base}}\, w_s(\tau)}{\sum_{s' \in \mathcal{S}} p_{s'}^{\text{base}}\, w_{s'}(\tau)}, \qquad \tau \in \{24, 36, 48\}\ \text{h}, \tag{6}$$

with $w_s(\tau) > 0$ monotone in intensity and inversely proportional to lead time: the numerator tilts each base probability based on the evidence available at τ, and the denominator restores a distribution. The whole pipeline is polynomial in the fixed-dimensional feature space.

Two-stage stochastic extension. $\mathrm{P}_{\mathrm{I}}^{\mathrm{S}}$ reuses the flow variables (x^F) as here-and-now decisions and adds recourse variables ($\delta_{s,i}^+, \delta_{s,i}^- \in \mathbb{Z}_+$) measuring persons left behind or evacuated in excess:

$$\min_{x^F, \delta^{\pm}} \sum_{(i,j) \in \mathcal{A}} c_{ij} \sum_{\text{id} \in \mathcal{I}} h_{\text{id}}\, x_{\text{id},i,j}^F + \sum_{s \in \mathcal{S}} p_s \sum_{i \in \mathcal{N}_A} \left(\theta^+ \delta_{s,i}^+ + \theta^- \delta_{s,i}^-\right), \tag{7a}$$

$$\text{s.t. constraints (2c) and (2d) of } \mathrm{P}_{\mathrm{I}}^{\mathrm{D}}, \qquad x_{\text{id},i,j}^F \in \mathbb{Z}_+,\ \delta_{s,i}^{\pm} \in \mathbb{Z}_+, \tag{7b}$$

$$\sum_{\text{id}: o_{\text{id}} = i} h_{\text{id}} \sum_{j:(i,j) \in \mathcal{A}} x_{\text{id},i,j}^F + \delta_{s,i}^+ - \delta_{s,i}^- = \sum_{h \in H} h f^s(h,i),\ \forall s,\ \forall i \in \mathcal{N}_A. \tag{7c}$$

The first term of (7a) is the cost of the plan committed before the storm resolves; the second is the expected cost of its mismatch with the realized demand. Equation (7c) states that the persons evacuated from i, plus those left behind, minus those moved in excess, equal the demand scenario (s) produced at i. The asymmetry of $\theta^+ \gg \theta^-$ encodes the ethical and operational cost of leaving people behind versus over-evacuating and pushes the first-stage plan towards protection against the tail scenario rather than nominal efficiency.

Static robust extension. For epistemic uncertainty on travel times, let $\mathcal{U}_{\text{box}} = \{\boldsymbol{\tau} : \tau_r \in [\bar{\tau}_r, \bar{\tau}_r(1+\rho^\tau)]\ \forall r\}$, with $\rho^\tau \ge 0$. Applying the worst-case substitution of Soyster [26], the duration constraint in (4c) is replaced by $T_{v,k}^{\text{fin}} \ge T_{v,k}^{\text{ini}} + \sum_r \tau_r^{\text{rob}} W_{v,r,k}$, with $\tau_r^{\text{rob}} = \lceil \bar{\tau}_r(1+\rho^\tau) \rceil$: every route is scheduled as if it took its longest admissible duration.

Proposition 1 (Robust feasibility and price of robustness)**.** *A solution feasible for* $\mathrm{P}_{\mathrm{II}}^{\mathrm{R}}$ *is feasible for every* $\boldsymbol{\tau} \in \mathcal{U}_{\text{box}}$*, since* $\tau_r \le \tau_r^{\text{rob}}$ *for all* r*; the decomposition is preserved because* $\mathcal{U}_{\text{box}}$ *is separable along* $r \in \Omega_o$*. Moreover,* $T^{\text{ms,rob}} \le (1+\rho^\tau) T^{\text{ms,nom}}$.

Affine adaptive extension. To reduce that conservatism, we adopt affine decision rules [25] over the budgeted set ($\mathcal{U}_\Gamma^{\mathcal{S}} = \{\boldsymbol{\xi} : \xi_r \in [-1,1], \sum_r |\xi_r| \le \Gamma\}$, where Γ limits how many routes may simultaneously deviate to their worst case; $\Gamma = |\Omega|$ recovers the box set, and $\Gamma = 0$ recovers the nominal model). With $T_{v,k}^{\text{fin}}(\boldsymbol{\xi}) = T_{v,k}^{0,\text{fin}} + \sum_r t_{v,k}^{r,\text{fin}} \sigma(r) \xi_r$, the schedule reacts linearly to the realized perturbation instead of being fixed against the worst case.

Theorem 2 (Tractable counterpart)**.** $T_{v,k}^{\text{fin}}(\boldsymbol{\xi}) \le T_{\max}$ *holds for all* $\boldsymbol{\xi} \in \mathcal{U}_\Gamma^{\mathcal{S}}$ *if and only if there exist* $\pi_{v,k} \ge 0$ *and* $\mu_{v,k,r}^{\pm} \ge 0$ *with*

$$T_{v,k}^{0,\text{fin}} + \Gamma \pi_{v,k} + \sum_{r \in \Omega} (\mu_{v,k,r}^+ + \mu_{v,k,r}^-) \le T_{\max}, \quad \pi_{v,k} + \mu_{v,k,r}^{\pm} \ge \pm\, \sigma(r)\, t_{v,k}^{r,\text{fin}}\ \ \forall r. \tag{8}$$

The worst case of the left-hand side is the value of an inner linear program with coefficients of $c_r = \sigma(r) t_{v,k}^{r,\text{fin}}$*, whose strong dual gives (8); the reformulation adds* $O(|\mathcal{V}||\mathcal{K}||\Omega|)$ *continuous variables and preserves the block structure, so Theorem 1 still applies. Consequently,* $Z_*^{\text{ARO}} \le Z_*^R$*, strictly whenever* $\Gamma < |\Omega|$ *and some* $\sigma(r) > 0$.

https://doi.org/10.3390/engproc2026153007

The deviation ($\sigma(r)$) is anchored to the ensemble ($\{(f^s, p_s^{\text{final}}(\tau))\}$)—it is proportional to the dispersion the ensemble induces on the duration of route r—rather than calibrated ad hoc as in classical budget-robust models [8]. This anchoring binds the stochastic and robust layers into a single hierarchy and addresses gap (G2).

6. Complexity, Solution Strategy and Operational Feasibility

6.1. Complexity of the Integrated Decision Problem

Definition 3 (Evac-Decision)**.** *Given a deterministic metamodel* ($\mathrm{P}_{\mathrm{I}}^{\mathrm{D}} \to \mathrm{P}_{\mathrm{II}}^{\mathrm{D}}$) *with network* $\mathcal{G}$, *partition* $\mathcal{N} = \mathcal{N}_A \sqcup \mathcal{N}_R \sqcup \mathcal{N}_F$, *demands* $\{f^0(h,i)\}$, *capacities* $\{u_{ij}, \alpha_i, \beta_j, \pi_j\}$, *costs* $\{c_{ij}, K_j\}$, *fleet* $\mathcal{V}$ *with* $\{\mathrm{Cap}_v, \mathrm{VehOrig}_{v,o}\}$, *horizon* $T_{\max}$, *trip bound* $K_{\max}$ *and threshold* κ *decide whether feasible solutions of* $\mathrm{P}_{\mathrm{I}}^{\mathrm{D}}$ *and* $\mathrm{P}_{\mathrm{II}}^{\mathrm{D}}$ *exist with a joint objective at most* κ *values.*

Theorem 3 (NP-hardness)**.** *Evac-Decision is NP-hard, even under topological segmentation, route consolidation and the decomposition of Theorem 1.*

Proof (sketch). Reduction from 3-SAT. Let φ be a 3-CNF formula on n variables and m clauses. A P_{II}*-trivial* configuration ($\mathcal{V} = \{v_1\}$, $\mathrm{Cap}_{v_1} = n$, $K_{\max} = 1$, $\mathrm{VehOrig}_{v_1,\cdot} = 1$, $\bar{\tau}_r = 0$, $T_{\max} = +\infty$) makes $\mathrm{P}_{\mathrm{II}}^{\mathrm{D}}$ trivially feasible, concentrating the difficulty on $\mathrm{P}_{\mathrm{I}}^{\mathrm{D}}$. *Variable gadget:* for each i, an origin (A_i) with $f^0(1, A_i) = 1$, transit nodes (T_i, F_i) and arcs ((A_i, T_i), (A_i, F_i)) of unit cost and unit capacity. *Clause gadget:* for each j, a shelter (S_j) with $\pi_{S_j} = +\infty$ and minimum coverage of $\sum_i x^F_{1,i,S_j} \geq 1$; for every literal ℓ_{jk}, add (T_i, S_j) if $\ell_{jk} = x_i$ or (F_i, S_j) if $\ell_{jk} = \neg x_i$, of zero cost and unbounded capacity. *Sink:* a global Sink $\in \mathcal{N}_F$ with $\pi_{\text{Sink}} = m$, zero-cost arcs (S_j, Sink), and threshold of $\kappa = n$. A satisfying assignment induces one unit-flow path per variable, with a total cost of $n = \kappa$, and each clause is satisfied iff some literal arc carries unit flow into S_j; conversely, any feasible solution with $z \leq n$ routes exactly one unit per variable by the unit arc capacities, defining a satisfying assignment. The construction has $O(n + m)$ nodes and arcs with parameters on $O(\log(n + m))$ bits. □

The extensions ($\mathrm{P}_{\mathrm{I}}^{\mathrm{S}}$, $\mathrm{P}_{\mathrm{II}}^{\mathrm{R}}$ and $\mathrm{P}_{\mathrm{II}}^{\mathrm{ARO}}$) inherit NP-hardness, each containing Evac-Decision under degenerate uncertainty. The decomposition, the polynomial route extraction, the polynomial MGEA preprocessing and the tractable counterpart of Theorem 2 are techniques to *cope with* this difficulty rather than remove it.

6.2. Exact Optimization with Structural Decomposition Rather than a Metaheuristic

Given that hardness, the natural expectation is a metaheuristic, as much of the literature reports [17,18]. We took the opposite decision for three reasons, the first being the nature of the objective: its dominant term penalizes uncovered demand, that is, people left behind, and without an optimality certificate, a planner cannot distinguish "the fleet is insufficient" from "the search did not find the assignment that would have sufficed"—for a plan handed to a Defense Council, the distinction is not academic. Second, the problem has an exploitable structure: authorization and the prefix definition of containment confine all coupling inside origin blocks, the master assignment is small, and each subproblem closes in tens of milliseconds; empirically, this collapses the binary space by 93.5% and keeps every reported solution provably optimal. We buy tractability from the structure of the problem rather than from the quality of the solution. The third reason is reproducibility: two runs of the same instance return the same plan with the same certificate, which matters when the model is re-run as the forecast updates. None of this argues against metaheuristics in general: they are the right tool for dense urban networks or continuous-time traffic dynamics, the setting of the works cited above. Beyond exact reach, the natural extension is to keep the exact master

https://doi.org/10.3390/engproc2026153007

and solve the per-origin subproblems heuristically, which is possible precisely because Theorem 1 isolates them.

6.3. Cost and Feasibility of the Sequential Pipeline

Solving five models in sequence raises a fair concern for emergency use; three properties mitigate it. *Only two layers are hard*: segmentation, capacity scaling, identifier encoding, route extraction and scenario generation are polynomial and, together, take less than a second; the two mixed-integer cores run in fractions of a second; and the stochastic and robust layers bring the total to about six minutes—roughly 320× the deterministic core and two orders of magnitude below the operational window (Table 3). *The pipeline degrades gracefully*: each layer yields a usable plan on its own, so the layers are cumulative refinements rather than mandatory stages, and if time runs out or a single track becomes credible, the planner stops at $P_I^D \rightarrow P_{II}^D$ with a plan in under a second. *Failure modes are legible*: an infeasibility is attributable either to network feasibility or to fleet adequacy, whereas a monolithic formulation would conflate the two into one message. The architecture does have a genuine cost, which we prefer to state: solving the layers in sequence is not equivalent to solving the joint problem, since the flow layer commits routes without anticipating the schedule that will serve them. That is the price of tractability, shared by every hierarchical planning scheme in transportation, and Section 9 indicates how a feedback loop could reduce it.

Table 3. Tractability boundary. Polynomial stages stay below 0.1 s; the two NP-hard cores and the inherited cascade dominate. Averages over the five instances.

Stage	Class	Time	Stage	Class	Time
Segmentation, capacity scaling	poly.	$<10^{-2}$ s	Flow layer P_I^D	NP-hard	0.128 s
MGEA, model selection, reweighting	poly.	$\sim 10^{-1}$ s	Scheduling P_{II}^D via DJO	NP-hard	0.077 s
Family-identifier encoding	poly.	$<10^{-2}$ s	Stochastic P_I^S	Inherited	∼20 s
Route extraction (Algorithm 1)	poly.	$<10^{-2}$ s	Robust P_{II}^R/P_{II}^{ARO}	Inherited	∼40/300 s

7. Results and Discussion

7.1. Case Study, Data Provenance and Setup

The framework is instantiated on Pinar del Río, the westernmost Cuban province, repeatedly affected by hurricanes entering the Gulf of Mexico. Origins, transit nodes and shelters are municipal centers, and arc distances follow the official road network; family exposure is disaggregated by size ($h \in \{1, \dots, 8\}$) from census data; the fleet is heterogeneous (buses with $\text{Cap}_v \approx 45$, minibuses and rural vans with 20–30); shelter capacities and activation costs follow the evacuation-center inventory of the provincial Defense Council, and the fraction of each municipality requiring shelter derives from the affectation probability reported by MGEA rather than assumed uniform. The event is hurricane **IAN** (Cat-4, September 2022), with a forecast horizon of $\tau = 12$ h before landfall and $H = 48$ h as the pre-impact window. Five instances use 6 origins with $|R| \in \{3, 4\}$, $|F| \in \{2, 3\}$, $|\mathcal{N}| \in \{11, 12, 13\}$; capacity is scaled by $\lambda_k = 1.10$, with $c = 0.36$ \$/km·person, $w_{\text{ms}} = 10$, $\rho^\tau \in \{0, \dots, 0.25\}$ and $\sigma_{\text{ARO}} = 0.15$. All models are solved with Google OR-Tools (CBC) on a commodity laptop (Intel i7, 16 GB), with timeouts of 30 s for the deterministic core and 120 s for the uncertainty layers; the implementation is openly available [11], so every number below is reproducible.

7.2. Deterministic Core, Stochastic Cascade and Tail Risk

Table 4 reports the deterministic core. All instances solve to optimality in both layers, with sub-second times (0.128 s for P_I^D, 0.077 s for P_{II}^D on average). Costs range from \$3.49 M

to \$5.26 M (CV = 15.3%) and makespans range from 22.5 to 45.0 h (mean of 34.2 h), inside the 48 h horizon; route extraction yields $|\Omega| \in [31, 55]$. The decomposition is engaged by default: binaries ($W_{v,r,k}$) fall from a monolithic baseline of about 22,000 to roughly 1600 per instance, corresponding to an empirical reduction of **93.5%**, consistent with Theorem 1(ii).

Two BIC-selected scenarios are sampled for the cascade: ω_1 (Cat-3, $p = 0.643$, $\eta = 1.00$) and ω_2 (Cat-4, $p = 0.357$, $\eta = 1.20$). The stochastic layer commits first-stage flows, absorbing demand uncertainty through $\delta^{\pm}$, with $E[Z^S] = \$26.9$ M, $\mathrm{CV}(Z^S) = 127.8\%$, $\mathrm{VaR}_{90} \approx \mathrm{CVaR}_{90} \approx \$71.3$ M and realized values of \$1.04 M and \$73.50 M. The dispersion is driven by the Cat-4 scenario, whose multiplier exceeds nominal shelter capacity in two segments; the ratio of $\mathrm{CVaR}_{90}/E[Z^S] \approx 2.65$ signals tail-driven cost and justifies the robust layers. On the sweep, the static makespan grows linearly with ρ^τ as Proposition 1 bounds, reaching +166.7% at $\rho^\tau = 0.25$, while the adaptive makespan grows at roughly $0.40(1 + \rho^\tau)$, recovering a conservatism gap near 16% of the nominal makespan.

Table 4. Deterministic experiment ($\mathrm{P}_{\mathrm{I}}^{\mathrm{D}} \rightarrow \mathrm{P}_{\mathrm{II}}^{\mathrm{D}}$) on the five IAN instances. Times in seconds, costs in U.S. dollars, and makespans in hours; OPT indicates proven optimality.

Inst.	$A \cdot R \cdot F$	$\lvert\mathcal{N}\rvert$	$\mathrm{P}_{\mathrm{I}}^{\mathrm{D}}$	Z^{D} (\$)	$t_{\mathrm{P}_{\mathrm{I}}^{\mathrm{D}}}$ (s)	$\lvert\Omega\rvert$	$\mathrm{P}_{\mathrm{II}}^{\mathrm{D}}$	T^{ms} (h)	$t_{\mathrm{P}_{\mathrm{II}}^{\mathrm{D}}}$ (s)
I_1	6·3·2	11	OPT	3,490,541	0.239	55	OPT	22.5	0.146
I_2	6·4·2	12	OPT	5,262,505	0.092	31	OPT	35.0	0.058
I_3	6·4·3	13	OPT	4,059,169	0.117	31	OPT	31.5	0.054
I_4	6·4·3	13	OPT	4,577,989	0.127	33	OPT	45.0	0.071
I_5	6·3·2	11	OPT	4,743,999	0.066	31	OPT	36.75	0.057
			Mean	4,426,841	0.128	36.2		34.15	0.077

7.3. Verification and Validation

Because the framework chains five models and four algorithms, we separate verification (Is the system solving the problem we intended to state?) from validation (Does that problem represent the evacuation it claims to represent?). Four levels of checks run automatically with every experiment and are part of the released code [11].

Model verification. For each solution, we recompute the flow-balance residual of (2b) per identifier and node, the slacks of (2c) and (2d) and the integrality of all integer variables; residuals must be exactly zero, and slacks must be non-negative; the battery aborts otherwise. The identity, i.e., $P_0 = \sum_{\mathrm{id}} h_{\mathrm{id}} q_{\mathrm{id}}$, is checked before optimization, catching inconsistencies between the census disaggregation and the encoded instance. Solver status is required to be OPT with zero gap, so the values in Table 4 are proven optima, not incumbents.

Algorithmic verification. Three properties of the solver are tested independently. Route extraction is verified by re-aggregation: the routes returned by Algorithm 1 must reproduce the optimal arc flows arc by arc, and their demands must sum to the total. The decomposition is verified by equivalence: on reduced instances that a monolithic formulation can still close, the master-plus-subproblems scheme attains the same objective value, which is the empirical counterpart of Theorem 1(i) and also the baseline against which the 93.5% reduction is measured. Robust feasibility is verified by Monte Carlo sampling of $\boldsymbol{\tau} \in \mathcal{U}_{\mathrm{box}}$, re-evaluating the schedule on each draw with no horizon violation admitted—the behavior Proposition 1 predicts, tested rather than assumed.

Consistency with the proven bounds. Measured behavior is confronted with the propositions it should obey: the robust makespan respects $T^{\mathrm{ms,rob}} \leq (1 + \rho^\tau) T^{\mathrm{ms,nom}}$ over the whole sweep, the adaptive model never exceeds the static one, and the per-origin binary count satisfies Theorem 1(ii) strictly in all instances. Agreement between proved bounds and measurements is strong verification evidence, since an error in the decomposition or in the robust counterpart would break it immediately.

Data and face validation. The instance is not synthetic: topology and distances come from the official provincial road network, shelter come capacities from the Defense Council registry, family-size disaggregation come from census data and the storm ensemble comes from HURDAT2. The reproduced event was managed under the same 48 h discipline the model assumes: the obtained makespans fall inside the window actually used by the authorities, and 88–91% shelter utilization is consistent with the prescribed reserve policy. We claim no more, since no controlled comparison against the realized evacuation is possible—the counterfactual plan was never executed.

7.4. Comparison with Recent State-of-the-Art Studies

Direct numerical benchmarking is not possible here, and it is better to say so than to present incommensurable figures: there is no shared benchmark library for hurricane evacuation; instances are built from the road network, census and shelter inventory of a specific territory; and objectives differ in kind—evacuation time [18,19], weighted travel and shelter cost [4], as well as expected multi-objective cost [24]. Structure, solution strategy and the behavior of the reported quantities can be compared.

On *uncertainty*, our results reproduce a pattern the stochastic literature reports consistently—namely, that expected cost is dominated by a low-probability, high-severity scenario: Heng et al. [24] documented the same asymmetry in flood scenarios, and Siddig and Song [23] motivated multi-stage formulations on this basis. Our ratio of $\mathrm{CVaR}_{90}/E[Z^S] \approx 2.65$ quantifies it. Where we depart is in the next step: instead of adding stages, we add a robust layer calibrated from the same ensemble, addressing a limitation those studies acknowledge but do not resolve, since Eriskin and Karatas [12] and the models of Section 2 fix the inflation exogenously, typically at 25%—which, in our sweep, expands the makespan by 166.7%, of which the adaptive model recovers about 16 percentage points. On *solution strategy*, the closest comparators are Bayram and Yaman [4], exact through Benders, and Khayamim et al. [19], whose decomposition-based ε-constraint algorithm generates Pareto fronts; both confirm that decomposition rather than approximation is what makes exact evacuation models viable. Ours differs in origin: Benders exploits a block structure after a partition chosen by the modeler, whereas our partition is dictated by the problem through vehicle–origin authorization, so the subproblems are independent rather than coupled by cuts. Against metaheuristic approaches [17,18], our runtimes are of the same order, with the difference being that solutions carry certificates. On *granularity*, Dulebenets et al. [18] and Khayamim et al. [19] incorporated vulnerable groups, which flow-based models rarely do, but the unit of assignment remains the individual; family indivisibility, as a commodity attribute, traceable down to the 185 family-level routes of the case study, is, to the best of our knowledge, specific to this work. On *complexity*, those works invoke hardness of the components they use, whereas Theorem 3 applies to the coupled problem actually solved and locates the boundary between hard cores and polynomial pipeline—which says which parts will degrade if the province grows and which will not.

8. From Model to Practice

The planner supplies three data objects that already exist in the civil-defense workflow: embarkation points with families disaggregated by size, prepared from the vulnerability studies and the projected track; the evacuation-center inventory with capacity, condition and activation cost; and the fleet, along with its authorization by municipality, with distances being read from the provincial network. No algorithmic parameter has to be tuned by the user: the penalties ($\theta^{\pm}$ and M_{μ}) and the weight (w_{ms}) encode policy, not solver behavior, and are set once with the responsible authority.

https://doi.org/10.3390/engproc2026153007

The output is not a vector of optimal values but three artifacts an operations room can act on: a shelter activation and family assignment map; a vehicle timetable in the form of Figure 2b, listing, for each vehicle, its trips, along with start and end times; and a small set of diagnostics—price of robustness, conservatism gap, tail-risk ratio and the coverage slack (μ_r, arguably the most useful in practice because it converts a diffuse worry ("do we have enough buses?") into a located, quantified deficit). The cadence follows the following forecast: the full chain at 48 h, re-run at 36 and 24 h as (6) reweights the ensemble at about six minutes per pass, and the deterministic core alone in the final hours, in under a second. The model schedules the pre-impact evacuation of families identified as vulnerable; it does not decide the evacuation order, manage return, or model self-evacuation by private vehicle, which is marginal in the Cuban setting but dominant elsewhere, where coupling with a traffic assignment model would be needed. Deployment was a design constraint: the implementation runs on a commodity laptop without commercial solver licenses [11], since a tool requiring a cluster or an expiring license is not usable by the institutions that would benefit most from it.

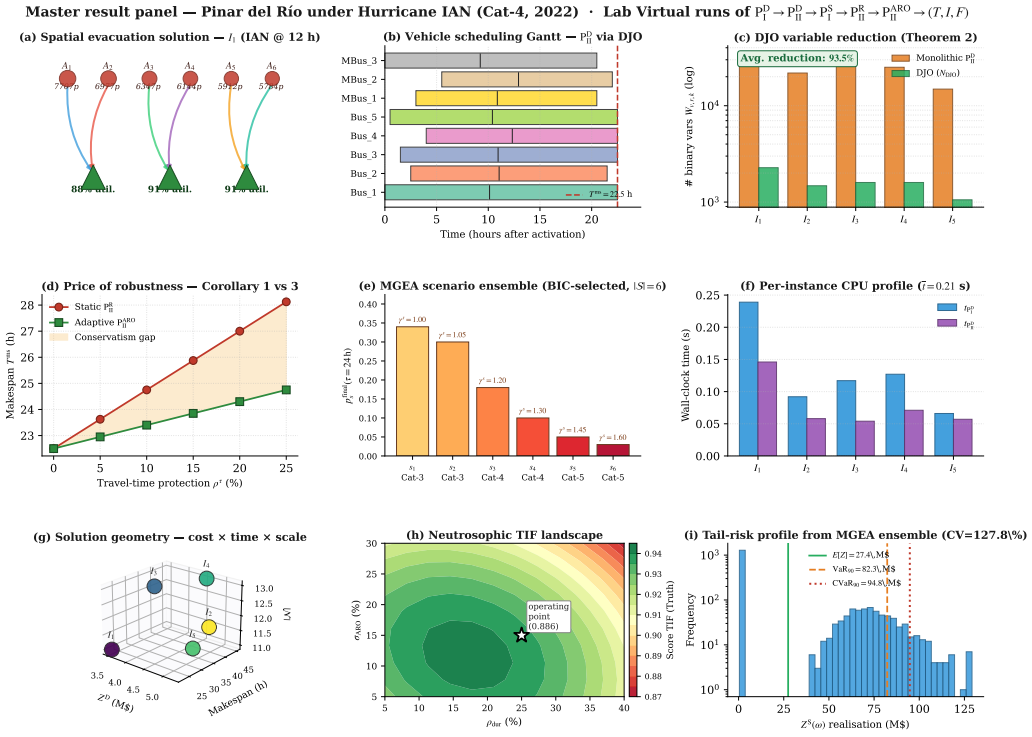


Figure 2. Master panel for Pinar del Río under hurricane IAN. (**a**) Routing of origins to active shelters by segment (88–91% utilisation). (**b**) Gantt chart of consolidated trips on I_1, $T^{\text{ms}} = 22.5$ h. (**c**) Decomposition reduction (mean 93.5%). (**d**) Price of robustness—static versus adaptive. (**e**) MGEA ensemble; $|\mathcal{S}| = 6$. (**f**) CPU profile. (**g**) Solution geometry. (**h**) Neutrosophic score landscape; TIF = 0.886. (**i**) Tail-risk profile of $Z^S(\omega)$.

9. Limitations and Future Research

No priority classes among evacuees. The model treats families of equal size as interchangeable, and this is its most consequential simplification: a plan optimal in makespan may evacuate a low-risk family before a high-risk one, and as it stands, the model would not object. Two features partially compensate. The demand set is not the general population but the families that the civil-defense vulnerability studies already classify as requiring assisted evacuation, so a coarse prioritization is applied upstream, and family indivisibility guarantees that children, elderly relatives and persons with reduced mobility are never separated from those who care for them, which aggregate formulations cannot ensure.

Neither substitutes for explicit priorities. The extension is tractable: replacing the uniform deficit penalty with an origin- and class-specific $\theta^{+}_{i,c}$ weighted by a vulnerability index makes the objective sensitive to *who* is left behind while preserving separability and therefore Theorem 1, and precedence constraints on trip indices so that higher-priority classes are served in earlier trips are equally compatible with the decomposition, as long as classes are defined within an origin. Ordering constraints coupling priorities across origins would break separability and require a different decomposition, which is why this deserves careful treatment rather than a footnote.

Other limitations. Sequential solution is not equivalent to joint optimization (Section 6.3); a feedback loop re-solving P^{D}_{I} with penalties reflecting the scheduling cost of its own routes would narrow the gap, and quantifying that gap on instances small enough for a monolithic solution is the natural next experiment. Results cover one province, one event and five instances, so extending the validation to the HURDAT2 catalog restricted to the years with reliable wind-radius data would strengthen generalization. Parameters other than demand and travel time—notably, shelter capacities, which can degrade during the event—are treated deterministically, for which distributionally robust formulations anchored to the same ensemble are the natural formalism. Fleet availability is assumed to be known, which matches the Cuban setting but not systems where vehicles must be requisitioned. Three further lines seem promising: neutrosophic calibration towards interval-valued sets, a re-optimization loop driven by real-time track updates, and equity objectives as soft constraints.

10. Conclusions

We presented a hierarchical metamodel for a priori family evacuation under hurricane threats, designed to close three gaps: the fragmentation between stochastic and robust treatments of uncertainty, the absence of a formal transformation from network flows to executable vehicle schedules, and the aggregation of evacuees into headcounts ignoring family indivisibility. The decomposition by origin (Theorem 1) yielded a 93.5% reduction in the binary decision space, with optimality preservation rather than approximation. Although the decision version is NP-hard by reduction from three-satisfiability, even under segmentation and consolidation (Theorem 3), the surrounding pipeline remains polynomial, concentrating the hard core in two well-defined layers. In the IAN case study, the deterministic core runs in sub-second time and completes the full robust cascade in about six minutes, with diagnostics—price of robustness of +166.7%, a conservatism gap near 16% recovered by the adaptive layer, $\mathrm{CVaR}_{90}/E[Z^{S}] \approx 2.65$ and a coverage slack that localizes fleet shortfalls municipality by municipality—that feed civil-defense decision-making directly. The main limitation, explicit priorities among evacuee classes, admits an extension compatible with the decomposition and heads the agenda, together with a feedback loop between layers, broader historical validation and distributionally robust formulations anchored to the same ensemble.

Author Contributions: Conceptualization, Y.F.-F. and S.M.A.-A.; methodology, formal analysis, investigation and original draft preparation, Y.F.-F.; review and editing, S.M.A.-A., R.M.-P. and G.B.-A.; software and visualization, K.L.O.C.; supervision, S.M.A.-A. and G.B.-A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The software implementation supporting the case study is openly available on Zenodo [11].

Acknowledgments: The authors thank the Faculty of Mathematics and Computation of the University of Havana, the Universidad Politécnica Estatal del Carchi (Ecuador) and the Universidad de Cienfuegos (Cuba), and the reviewers, whose observations led to a substantially clearer version of this work.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Hamacher, H.W.; Tjandra, S.A. Mathematical modelling of evacuation problems: A state of the art. In *Pedestrian and Evacuation Dynamics*; Schreckenberg, M., Sharma, S.D., Eds.; Springer: Berlin/Heidelberg, Germany, 2002; pp. 227–266.
2. Murray-Tuite, P.; Wolshon, B. Evacuation transportation modeling: An overview of research, development, and practice. *Transp. Res. Part C Emerg. Technol.* **2013**, *27*, 25–45. [CrossRef]
3. Yin, Y.; Xu, X.; Wang, D.; Yu, Y.; Cheng, T.C.E . Two-stage recoverable robust optimization for an integrated location-allocation and evacuation planning problem. *Transp. Res. B* **2024**, *182*, 102906. [CrossRef]
4. Bayram, V.; Yaman, H. Shelter location and evacuation route assignment under uncertainty: A Benders decomposition approach. *Transp. Sci.* **2018**, *52*, 416–436. [CrossRef]
5. Birge, J.R.; Louveaux, F. *Introduction to Stochastic Programming*, 2nd ed.; Springer: New York, NY, USA, 2011.
6. Kall, P.; Wallace, S.W. *Stochastic Programming*; Wiley: Chichester, UK, 1994.
7. Ben-Tal, A.; Nemirovski, A. Robust solutions of uncertain linear programs. *Oper. Res. Lett.* **2000**, *25*, 1–13. [CrossRef]
8. Bertsimas, D.; Sim, M. The price of robustness. *Oper. Res.* **2004**, *52*, 35–53. [CrossRef]
9. Ben-Tal, A.; El Ghaoui, L.; Nemirovski, A. *Robust Optimization*; Princeton University Press: Princeton, NJ, USA, 2009.
10. Fernández-Fernández, Y.; Allende-Alonso, S.; Miranda-Pérez, R.; Bouza-Allende, G.; Narváez-Quendi, J. Scalable network flow optimization for natural hazards emergency evacuation: Safe path for socially vulnerable communities. In *Emerging Research in Intelligent Systems. CIT 2025*; Lecture Notes in Networks and Systems; Springer: Cham, Switzerland, 2026; Volume 1881, pp. 71–86.
11. Fernández-Fernández, Y. *Metamodelo de Optimización Robusta para Planificación de Evacuaciones Ante Huracanes*, version 1.0.0 [Software]; Zenodo: Geneva, Switzerland, 2026. [CrossRef]
12. Eriskin, L.; Karatas, M. A novel robust optimization model for hurricane evacuation under uncertainty. *Omega* **2024**, *122*, 102959.
13. Kitchenham, B.; Charters, S. *Guidelines for Performing Systematic Literature Reviews in Software Engineering*; EBSE Technical Report EBSE-2007-01; Keele University and University of Durham: Keele, UK, 2007.
14. Esposito Amideo, A.; Scaparra, M.P.; Kotiadis, K. Optimising shelter location and evacuation routing operations: The critical issues. *Eur. J. Oper. Res.* **2019**, *279*, 279–295. [CrossRef]
15. Swamy, R.; Kang, J.E.; Batta, R.; Chung, Y. Hurricane evacuation planning using public transportation. *Socio-Econ. Plan. Sci.* **2017**, *59*, 43–55. [CrossRef]
16. Zhao, X.; Ji, K.; Xu, P.; Qian, W.; Ren, G.; Shan, X. A round-trip bus evacuation model with scheduling and routing planning. *Transp. Res. A* **2020**, *137*, 285–300. [CrossRef]
17. Goerigk, M.; Deghdak, K.; Heßler, P. A comprehensive evacuation planning model and genetic solution algorithm. *Transp. Res. E* **2015**, *63*, 124–139.
18. Dulebenets, M.A.; Pasha, J.; Abioye, O.F.; Kavoosi, M.; Ozguven, E.E.; Moses, R.; Boot, W.R.; Sando, T. Exact and heuristic solution algorithms for efficient emergency evacuation in areas with vulnerable populations. *Int. J. Disaster Risk Reduct.* **2019**, *39*, 101114. [CrossRef]
19. Khayamim, R.; Moses, R.; Ozguven, E.E.; Borowska-Stefańska, M.; Wiśniewski, S.; Ozkul, S.; Dulebenets, M.A. A bi-objective optimization model and solution approach for emergency evacuation planning under pandemic settings. *Heliyon* **2025**, *11*, e43633. [CrossRef]
20. Rawls, C.G.; Turnquist, M.A. Pre-positioning of emergency supplies for disaster response. *Transp. Res. B* **2010**, *44*, 521–534. [CrossRef]
21. Vahdani, B.; Veysmoradi, D.; Noori, F.; Mansour, F. Two-stage multi-objective location–routing–inventory model for humanitarian logistics network design under uncertainty. *Int. J. Disaster Risk Reduct.* **2018**, *27*, 290–306. [CrossRef]
22. Karabuk, S.; Manzour, H. A multi-stage stochastic program for evacuation management under tornado track uncertainty. *Transp. Res. E* **2019**, *124*, 128–151. [CrossRef]
23. Siddig, M.; Song, Y. Multi-stage stochastic programming for hurricane evacuation under demand uncertainty. *Manag. Sci.* **2025**, *in press*. [CrossRef]
24. Heng, S.; Phutthanawong, K.; Boonmee, C. Optimizing localized humanitarian logistics: A stochastic programming approach for facility location, inventory and evacuation strategies. *J. Humanit. Logist. Supply Chain Manag.* **2025**, *ahead of print*. [CrossRef]
25. Ben-Tal, A.; Goryashko, A.; Guslitzer, E.; Nemirovski, A. Adjustable robust solutions of uncertain linear programs. *Math. Program.* **2004**, *99*, 351–376. [CrossRef]

26. Soyster, A.L. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Oper. Res.* **1973**, *21*, 1154–1157. [CrossRef]
27. Lenstra, J.K.; Rinnooy Kan, A.H.G. Complexity of vehicle routing and scheduling problems. *Networks* **1981**, *11*, 221–227. [CrossRef]
28. Garey, M.R.; Johnson, D.S. *Computers and Intractability: A Guide to the Theory of NP-Completeness*; W.H. Freeman: New York, NY, USA, 1979.
29. Fernández-Fernández, Y.; Allende-Alonso, S.M.; Miranda-Pérez, R.; Bouza-Allende, G. Two-stage evacuation model considering uncertainty of storm tracks in the Caribbean. *Int. J. Appl. Eng. Technol.* **2023**, *5*, 345–362.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.