



Application of a Heuristic-Based System for Sustainable Faculty Routing and Workload Balancing in Cooperative Education

Aritath Siraphatwongkorn¹, and Jirayus Arbking^{*}

¹ Faculty of Informatics, Burapha University, Chonburi, 20131, Thailand

^{*} Correspondence: jirayu.ar@go.buu.ac.th

Citation:

Siraphatwongkorn, A.; Arbking, J. Application of a heuristic-based system for sustainable faculty routing and workload balancing in cooperative education. *ASEAN J. Sci. Tech. Report.* **2026**, 29(9), e262669. <https://doi.org/10.55164/ajstr.v29i9.262669>

Article history:

Received: December 4, 2025

Revised: April 24, 2026

Accepted: May 16, 2026

Available online: July 26, 2026

Publisher's Note:

This article is published and distributed under the terms of the Thaksin University.

Abstract: In cooperative education programs, faculty supervisors are required to visit multiple industry sites hosting students, with each site imposing a distinct workload characterized by the number of assigned students, travel distance, and on-site activities. This study formulates the supervisory routing task as a Set Partitioning Vehicle Routing Problem with Equitable Workload Constraints (SPVRP-EW) and introduces a three-phase heuristic approach. The first phase generates feasible one-day trips within a strict time window constraint. The second phase selects a minimum set of trips that covers all locations using a greedy set-cover strategy with an efficiency-based tie-breaking rule. The final phase assigns the selected trips to six faculty supervisors to minimize the maximum–minimum workload gap while ensuring that each supervisor is responsible for an equal number of main trips. Using a real dataset of 24 placement sites across four provinces in Thailand, the proposed method generated 12 trips totaling 2,399.5 km. The solution achieved a maximum–minimum workload gap of only 2 points and a workload variance of 0.58. Extended experiments further demonstrate that the method outperforms two simulated manual plans, as well as random-cover and random-assignment baselines, in terms of total distance, number of trips, and workload variance. The heuristic remains robust under ± 5 –10% distance noise, with the workload gap staying within 4 points. In addition, synthetic instances with 50 and 100 locations can be solved in under one second, confirming their practical scalability. Relaxing the equity threshold beyond 2 points does not yield further improvement, indicating that the heuristic is already optimal for the studied instance. Overall, the framework serves as a decision-support tool that enhances both logistical efficiency and social sustainability by improving workload fairness in cooperative education supervision.

Keywords: Vehicle routing problem; workload equity; set cover heuristic; cooperative education; sustainability

1. Introduction

The Bachelor of Science Program in Software Engineering at the Faculty of Informatics, Burapha University, prioritizes industry-based learning to ensure that fourth-year students attain workplace readiness through the Cooperative Education (Co-op) program [1–3]. Each year, faculty members must travel to multiple host organizations to supervise students, assess learning outcomes, and maintain industry relationships. As cohorts grow and partner firms disperse

geographically, supervision becomes a substantial operational burden, with its fairness having direct consequences for faculty well-being and program sustainability [4]. Operational complexity is acute. Partner organizations are distributed across Bangkok, Nonthaburi, Pathumthani, and Chonburi provinces, resulting in heterogeneous travel times and workloads. Each supervisory day is bounded by a 12-hour operational ceiling (06:00–18:00). Every site requires an on-site activity whose duration varies by program structure, and a strict institutional equity policy requires that the difference between the heaviest and lightest workload carried by any supervisor is minimal.

Viewed through the lens of operations research, the problem resembles a constrained Vehicle Routing Problem (VRP) with time windows [5–7] and an equitable load-balancing requirement. Classical VRP formulations and their load-balancing variants have been studied extensively [8–12]; however, existing load-balancing models typically minimize makespan or tour length deviation across homogeneous vehicles, whereas cooperative education supervision imposes (1) categorical point-based workload rather than continuous load, (2) an equal number of main-responsible trips per supervisor, and (3) a hard equity threshold derived from institutional policy. To the best of our knowledge, no integrated formulation that simultaneously handles set-partitioning coverage and discrete workload equity in an academic context has been reported.

Consequently, this study develops a three-phase heuristic tailored to Burapha University's cooperative education model. The system (1) enumerates feasible daily routes within the 12-hour window, (2) selects a minimum number of trips that collectively cover every placement site using a greedy set-cover rule, and (3) assigns trips to faculty to minimize the maximum–minimum workload gap subject to an identical number of main-responsible trips per person. The contributions are: (1) a concrete formulation of the Set Partitioning VRP with Equitable Workload constraints; (2) an efficient decomposition with polynomial-time greedy selection and bounded branch-and-bound assignment; (3) empirical validation against multiple baselines, sensitivity analyses, scalability tests up to 100 locations, and robustness tests under distance noise; and (4) a sustainability analysis linking routing efficiency with social and environmental proxies.

2. Materials and Methods

2.1 Problem Definition, Data Input, and Constraints

2.1.1 Problem Context and Objective

The core problem is to plan supervisory routes for students placed at 24 locations (P_2 – P_{24}) originating from the Supervision Base P_1 . The model must ensure that every location is visited across multiple days, each day forming one round trip that starts and ends at P_1 , and that the resulting workload points are distributed equitably among six faculty supervisors, $F = \{A, B, C, D, E, F\}$. Formally, the decision problem is to determine a minimum-cardinality set of feasible daily trips that collectively cover all placement sites, along with an assignment of these trips to supervisors that minimizes the maximum–minimum gap in main-responsible workload points.

2.1.2 Data Input and Movement Parameters

The model uses the following defined parameters and data inputs (Table 1).

Table 1. Model Parameters, Descriptions, and Roles

Parameter	Description	Role
$D = [d_{ij}]$	24×24 symmetric driving-distance matrix (km)	Input to Phase 1 feasibility check
v	Constant driving speed, 60 km/h	Converts distance to travel time
vt_i	On-site activity (visiting) time at location i (min)	Component of total trip duration
p_i	Discrete workload points are attached to location i , proportional to the number of students and assessment depth.	Input to Phase 3 workload balancing
$[T_{min}, T_{max}]$	Admissible trip duration window [400, 720] min	Feasibility constraint for every trip
$F = \{A, \dots, F\}$	Set of six faculty supervisors	Assignment targets in Phase 3

2.1.3 Personnel Roles and Constraints

The assignment involves six faculty members. The operational rules are:

- Trip Crew Structure: Each trip is executed by two faculty members — a Main Responsible and a Partner.
- Point Attribution Rule: Only the Main Responsible receives the total workload points p_i accumulated at the locations visited during the trip.
- Time Window Constraint: Each trip starts at 06:00 and must end by 18:00, giving a maximum trip duration $T_{\max} = 720$ min. A minimum working duration $T_{\min} = 400$ min avoids under-utilized one-day plans (except for the administrative self-visit at P_1).
- Equal Main-Trip Constraint: Each supervisor must act as the Main Responsible for exactly $|K|/|F|$ trips; this is enforced explicitly during the Phase-3 assignment.
- Equity Constraint: The final assignment must satisfy the institutional fairness condition $\max_f W_f - \min_f W_f \leq 3$, where W_f is the total workload accumulated by faculty member f .

2.2 Phase 1: Candidate Trip Generation

Phase 1 enumerates feasible round trips. A trip k is defined as $k = (P_1, v_1, v_2, \dots, v_m, P_1)$, where $v_1, \dots, v_m \in \{P_2, \dots, P_{24}\}$. The total duration $T(k)$ is calculated by equation (1):

$$\sum_{i=0}^{m-1} \left(\frac{d_{v_i, v_{i+1}}}{v} \right) \times 60 + \sum_{i=1}^m v t_{v_i} \quad (1)$$

An equation determines feasibility. (2):

$$T_{\min} \leq T(k) \leq T_{\max} \quad (2)$$

Only trips satisfying equation (2) are retained in the candidate pool C . To limit combinatorial growth, trips of length > 3 are excluded a priori because any four-stop route violates the 720-minute ceiling on the studied dataset (the shortest feasible four-stop tour exceeds 740 min). For each feasible multi-stop combination, all permutations are retained so that Phase 2 may select the most efficient ordering.

2.3 Phase 2: Greedy Set Cover Selection

A greedy heuristic [13] is employed to select a minimum number of trips K from C such that every location is covered at least once. At each iteration, the heuristic chooses the trip that maximizes the number of newly covered locations (equation (3)) and, in the case of a tie, the trip with the shortest duration:

$$k^* = \arg \max_{k \in C} |V(k) \cap U| \quad (3)$$

where $V(k)$ is the set of locations visited by trip k , and U is the set of locations not yet covered. This rule corresponds to the classical Chvátal–Johnson greedy heuristic for minimum set cover, which delivers an $H(n) \approx \ln n + 1$ approximation to the optimum [13, 14]. The rule is selected because: (1) the placement-site instance is of small-to-medium size ($|C| < 10^3$). Exact integer programming would be tractable but not needed; (2) the greedy rule is deterministic, interpretable to program administrators, and implementable without a commercial solver; and (3) empirical tests (Section 3.7) show that the greedy policy yields almost 44% shorter total distance and 31% fewer trips than a random feasible cover.

2.4 Phase 3: Equitable Assignment and Output

Phase 3 assigns the selected trips K to the six faculty members in F . The workload $W(k)$ The cost of each trip is defined by equation (4):

$$W(k) = \sum_{i \in V(k)} p_i \quad (4)$$

Let $x_{k,f} = 1$ if faculty member f is Main Responsible for trip k and 0 otherwise. The goal is to find an assignment X that minimises the difference between the maximum and minimum total workload received by any faculty member, as shown in the equation. (5):

$$\min \left(\max_{f \in F} \sum_{k \in K} x_{k,f} W(k) - \min_{f \in F} \sum_{k \in K} x_{k,f} W(k) \right) \quad (5)$$

subject to $\sum_f x_{k,f} = 1; \forall k$ and $\sum_k x_{k,f} = |K|/|F|; \forall f$. The first constraint enforces one Main Responsible per trip, while the second enforces an equal number of main-responsible trips per supervisor. For $|K| \leq 14$ the assignment is solved exactly by a bounded branch-and-bound (B&B) procedure [15] that prunes when the current partial gap exceeds the best known. For larger instances, a Longest-Processing-Time (LPT)-style greedy fallback [16] provides a solution in $O(|K| \log |K|)$ time.

2.5 Comparison with Existing VRP Load Balancing Approaches

Load-balancing VRP variants aim to equalize a continuous attribute, usually tour length, duration, or demand across homogeneous vehicles [8, 9, 10]. Representative formulations include the Min–Max VRP [17], the balanced CVRP [18], and recent learning-based balanced VRPs [19]. Table 2 compares the present formulation (Set Partitioning VRP with Equitable Workload, SPVRP-EW) with these established variants.

Table 2. Comparison of the proposed SPVRP-EW with representative VRP load-balancing formulations.

Aspect	Min–Max VRP [15]	Balanced CVRP [18]	Learning-Based BVRP [19]	SPVRP-EW (Proposed)
Balancing target	Tour length	Vehicle demand	Tour duration	Discrete workload points
Balance metric	max route length	deviation from the mean load	route-duration variance	$\max - \min \text{ workload} \leq \tau$
Routing decision	Route per vehicle	Route per vehicle	Route per vehicle	Set of feasible one-day trips
Equity threshold	Implicit	Soft penalty	Soft penalty	Hard policy bound (institutional)
Trips per agent	Not enforced	Not enforced	Not enforced	Equal ($ K / F $)
Interpretability	Medium	Medium	Low (neural)	High (greedy + B&B)
Typical size	Small–medium	Medium–large	Large	Small–medium (≤ 100)

The comparison highlights three distinctive gaps addressed by the proposed SPVRP-EW. First, the workload unit is a discrete point rather than a continuous length, which reflects real academic assessment effort and precludes metric-based balancing. Second, equity is enforced as a hard max – min threshold rather than as a soft penalty, which aligns with the institutional fairness policy. Third, an equal-number-of-main-trips constraint is introduced, which is uncommon in logistics but essential in academic contexts to prevent pipeline inequality. Because the instance class is small to medium and the users are program administrators with limited access to solver infrastructure, a greedy set-cover plus branch-and-bound assignment is preferred over metaheuristics or exact mixed-integer programming for computational efficiency, interpretability, and maintainability.

2.6 Rationale of Experimental Design

The 24 placement sites are the actual locations that hosted Software Engineering cooperative-education students during the 2023–2024 academic year; hence, 24 is neither arbitrary nor artificially convenient. This dataset represents the typical operational scale of a single program at a Thai faculty and is the scale at which automated support is most needed. The equity threshold $\max - \min \leq 3$ is derived from the institution’s internal faculty-load regulation, which tolerates at most three units of disparity before a review is triggered; choosing this threshold therefore aligns the optimization target with an operational rule rather than a statistical convention. To stress-test both the method and the policy, Section 3.4 sweeps the equity threshold $\tau \in \{1, 2, 3, 4\}$, while Section 3.5 scales the instance up to $n = 100$ synthetic sites.

2.7 Computational Complexity Analysis

Let n be the number of placement sites and $|C|$ the size of the candidate pool. Phase 1 enumerates all single, ordered pairs, and ordered triples of short-visit locations, giving $|C| = O(n + n^2 + n^3)$. Each

duration check is $O(1)$, so Phase 1 is $O(n^3)$ in the worst case, which remains tractable for $n \leq 100$ on a standard workstation (Table 8, below). Phase 2 is the classical greedy set cover with runtime $O(|C| \cdot n)$, and it inherits the $H(n) = \sum_{i=1}^n 1/i \approx \ln n + 1$ approximation ratio [13, 14]. Phase 3, when solved exactly, is $O(|F|^{|K|})$ in the worst case, but (i) the equal-main-count constraint prunes the branch-and-bound tree and (ii) an upper bound on the running gap, which in practice finishes in a fraction of a second for $|K| \leq 14$. The LPT fallback is $O(|K| \log |K|)$. Overall, the end-to-end heuristic completes in sub-second time on the studied instance and in under 1 second for instances with up to 100 sites.

3. Results and Discussion

3.1 Placement Site Data and Distance Matrix

Table 3. Placement site data (visiting time and workload points).

Location i	Visit Time vt_i (min)	Points p_i	Location i	Visit Time vt_i (min)	Points p_i
P_1	62	2	P_{13}	110	5
P_2	62	2	P_{14}	94	4
P_3	94	4	P_{15}	110	5
P_4	62	2	P_{16}	94	4
P_5	62	2	P_{17}	158	8
P_6	62	2	P_{18}	110	5
P_7	46	1	P_{19}	62	2
P_8	174	9	P_{20}	206	11
P_9	158	8	P_{21}	46	1
P_{10}	78	3	P_{22}	46	1
P_{11}	62	2	P_{23}	46	1
P_{12}	94	4	P_{24}	46	1

Table 3 summarizes the 24 placement sites, their visiting times vt_i , and workload points p_i . Figure 1 depicts the spatial distribution across Bangkok, Nonthaburi, Pathumthani, and Chonburi provinces. The full 24×24 symmetric driving-distance matrix D . For example, the driving distance from P_1 (Supervision Base) to P_{24} is 129.89 km, which influences whether a trip containing P_{24} satisfies the 720-minute ceiling.




Figure 1. Spatial Distribution of the 24 Participating Establishments Across Four Provinces

Table 4. Distance Matrix (km)

From	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	P13	P14	P15	P16	P17	P18	P19	P20	P21	P22	P23	P24
P1	0.0	0.0	0.0	0.4	0.6	0.4	18.9	2.3	21.0	24.8	92.0	28.0	4.9	6.9	89.1	103.6	80.1	118.4	97.0	104.6	111.7	109.1	99.4	129.9
P2	0.0	0.0	0.0	0.4	0.6	0.4	18.9	2.3	21.0	24.8	92.0	28.0	4.9	6.9	89.1	103.6	80.1	118.4	97.0	104.6	111.7	109.1	99.4	129.9
P3	0.0	0.0	0.0	0.4	0.6	0.4	18.9	2.3	21.0	24.8	92.0	28.0	4.9	6.9	89.1	103.6	80.1	118.4	97.0	104.6	111.7	109.1	99.4	129.9
P4	0.4	0.4	0.4	0.0	0.3	0.7	19.2	2.6	21.3	22.6	92.3	28.3	5.2	7.2	89.4	103.9	80.4	118.7	97.3	104.8	112.0	109.3	99.7	130.2
P5	0.6	0.6	0.6	0.3	0.0	1.0	19.4	2.8	21.5	22.5	92.5	28.4	5.4	7.4	89.6	104.0	80.6	118.9	97.4	105.0	112.2	109.5	99.9	130.4
P6	0.4	0.4	0.4	0.7	0.9	0.0	19.0	4.1	21.1	24.9	92.1	28.1	4.4	7.0	89.2	103.6	80.2	118.5	97.0	104.6	111.8	109.1	99.5	130.0
P7	17.7	17.7	17.7	17.6	17.8	17.4	0.0	18.0	23.7	29.5	78.9	11.3	18.3	14.9	72.8	94.2	67.0	105.3	85.7	91.4	98.6	95.9	90.0	116.8
P8	3.7	3.7	3.7	3.7	3.8	3.4	19.0	0.0	20.0	26.3	92.1	28.1	2.0	7.1	89.3	103.7	80.2	118.6	97.1	104.7	111.8	109.2	99.5	130.0
P9	21.4	21.4	21.4	21.4	21.5	21.2	25.5	20.8	0.0	17.2	107.8	33.8	21.2	17.0	94.2	115.6	90.3	131.8	107.1	117.9	115.0	122.4	111.4	143.2
P10	26.5	26.5	26.5	24.2	24.1	26.2	34.0	28.4	21.4	0.0	116.3	42.3	28.8	28.4	102.7	124.1	98.8	140.2	115.6	126.4	123.5	130.9	119.9	151.7
P11	93.1	93.1	93.1	93.1	93.2	92.8	83.3	93.4	104.6	110.3	0.0	83.4	93.8	90.4	22.6	8.6	13.0	29.9	10.0	16.0	23.1	20.5	8.4	41.3
P12	28.2	28.2	28.2	28.2	28.3	27.9	10.8	28.5	33.2	39.0	79.7	0.0	28.9	25.5	72.2	93.6	67.8	109.8	85.2	95.9	93.0	100.4	89.5	121.2
P13	3.5	3.5	3.5	3.5	3.6	3.2	17.1	1.7	18.0	24.4	90.2	26.2	0.0	5.1	87.3	101.8	78.3	116.6	95.2	102.7	109.9	107.3	97.6	128.1
P14	5.8	5.8	5.8	5.8	5.9	5.6	19.4	6.1	16.4	22.2	92.5	28.5	6.5	0.0	87.1	108.4	80.6	118.9	100.0	105.0	112.2	109.6	104.3	130.4
P15	87.5	87.5	87.5	87.5	87.6	87.2	73.1	87.8	94.2	100.0	24.2	73.0	88.2	84.8	0.0	31.9	21.5	33.0	18.5	25.4	27.2	29.9	27.8	49.1
P16	107.8	107.8	107.8	107.7	107.9	107.5	90.4	108.1	111.6	117.4	9.1	90.4	108.5	105.1	29.6	0.0	22.4	25.5	9.1	8.0	15.8	9.2	0.9	37.0
P17	79.8	79.8	79.8	79.8	79.9	79.6	65.4	80.1	88.9	94.6	13.4	63.6	80.5	77.1	22.8	24.5	0.0	41.4	17.6	27.5	34.6	32.0	20.3	52.8
P18	112.8	112.8	112.8	112.8	112.9	112.5	98.4	113.1	119.5	125.3	30.0	98.3	113.5	110.1	31.2	24.2	39.2	0.0	22.4	21.3	14.6	24.6	24.6	16.4
P19	99.5	99.5	99.5	99.5	99.6	99.3	85.1	99.8	107.3	113.1	11.7	86.1	100.2	96.8	19.3	9.2	21.0	23.7	0.0	9.8	17.0	14.3	8.4	35.1
P20	111.7	111.7	111.7	111.7	111.8	111.5	94.3	112.0	115.6	121.3	14.4	94.4	112.4	109.0	30.2	7.1	26.4	23.3	10.0	0.0	10.4	6.4	7.5	31.3
P21	107.8	107.8	107.8	107.7	107.9	107.5	93.3	108.1	114.5	120.3	22.6	93.3	108.4	105.0	26.2	15.3	34.5	19.3	18.3	9.9	0.0	14.2	15.7	24.2
P22	118.1	118.1	118.1	118.0	118.2	117.8	100.7	118.3	121.9	127.7	20.7	100.7	118.7	115.3	33.2	8.1	32.7	26.3	14.5	7.0	13.9	0.0	9.0	37.8
P23	107.5	107.5	107.5	107.4	107.6	107.2	90.1	107.8	111.3	117.1	8.8	90.1	108.2	104.8	29.3	0.8	22.1	25.9	8.8	8.4	16.1	10.0	0.0	37.3
P24	125.3	125.3	125.3	125.3	125.4	125.0	110.9	125.6	132.0	137.8	43.6	110.8	126.0	122.6	43.7	37.7	52.8	17.4	35.7	32.8	26.1	38.2	38.2	0.0

Table 4 presents the symmetrical Distance Matrix D in kilometers used as a primary input for the heuristic model. This confirms that P_1 , P_2 , and P_3 all represent the same geographical location—the Faculty of Informatics building at Burapha University—which serves as the central Supervision Base.

3.2 Proposed Heuristic Outcome

Table 5. Final trip schedule and workload assignment.

#	Route	Total (min)	Drive (min)	Visit (min)	Main	Partner	Points
1	$P_1 \rightarrow P_4 \rightarrow P_{19} \rightarrow P_{22} \rightarrow P_1$	400.0	230.0	170.0	A	B	5
2	$P_1 \rightarrow P_{21} \rightarrow P_{24} \rightarrow P_7 \rightarrow P_1$	402.4	264.4	138.0	B	C	3
3	$P_1 \rightarrow P_{23} \rightarrow P_{10} \rightarrow P_5 \rightarrow P_1$	427.2	241.2	186.0	C	D	6
4	$P_1 \rightarrow P_2 \rightarrow P_{11} \rightarrow P_{22} \rightarrow P_1$	400.6	230.6	170.0	D	E	5
5	$P_1 \rightarrow P_{15} \rightarrow P_{13} \rightarrow P_1$	400.8	180.8	220.0	A	F	10
6	$P_1 \rightarrow P_{16} \rightarrow P_{14} \rightarrow P_1$	402.5	214.5	188.0	E	A	8
7	$P_1 \rightarrow P_{18} \rightarrow P_6 \rightarrow P_1$	403.3	231.3	172.0	E	B	7
8	$P_1 \rightarrow P_3 \rightarrow P_{17} \rightarrow P_1$	411.9	159.9	252.0	B	D	12
9	$P_1 \rightarrow P_{20} \rightarrow P_{12} \rightarrow P_1$	527.1	227.1	300.0	F	D	15
10	$P_1 \rightarrow P_1 \rightarrow P_1$	62.0	0.0	62.0	F	E	2
11	$P_1 \rightarrow P_8 \rightarrow P_{11} \rightarrow P_1$	423.5	187.5	236.0	D	A	11
12	$P_1 \rightarrow P_{23} \rightarrow P_9 \rightarrow P_1$	436.2	232.2	204.0	C	B	9

Final workload per supervisor: $W_A = 15$; $W_B = 15$; $W_C = 15$; $W_D = 16$; $W_E = 15$; $W_F = 17$, max – min = 2.

Applying the three-phase heuristic to the 24-site instance, Phase 1 produced. $|C| = 861$ feasible candidate trips, Phase 2 selected $|K| = 12$ trips that cover every location, and Phase 3 assigned those trips to six faculty members with a maximum–minimum workload gap of $\max_f W_f - \min_f W_f = 2$ points and a workload variance of 0.58. The selected trip schedule is reported in Table 5. All trips satisfy the [400, 720] min duration window. The total driving distance is 2,399.5 km. Final workloads are $W_A = 15$, $W_B = 15$, $W_C = 15$, $W_D = 16$, $W_E = 15$, $W_F = 17$ points, each member being the Main Responsible on exactly $|K|/|F| = 2$ trips, fully satisfying the equal-main-trip constraint.

3.3 Baseline Comparison

Table 6. Performance comparison across six methods on the 24-site instance.

Method	Trips	Distance (km)	max – min	Variance	Runtime (s)
Manual – Proximity Zones (2024)	14	~2,831.4	4	6.67	Several hours
Manual – Load-Balanced (Expert)	13	~2,615.5	3	4.17	Several hours
Random Cover + Random Assignment	17.33 (avg)	4,253.9 (avg)	41.00 (avg)	—	< 0.01
Greedy Cover + Random Assignment	12	2,399.5	32.00 (avg, max 57)	—	< 0.01
Simple Greedy (Greedy Cover + LPT)	12	2,399.5	2 *	0.58	< 0.01
Proposed (Greedy + Equity B&B)	12	2,399.5	2	0.58	0.27

* Simple Greedy LPT coincides with the optimum on this instance but provides no guarantee on equal-main-trip compliance or on the hard equity bound $\max - \min \leq \tau$; the proposed method provides both guarantees.

The proposed heuristic is compared with five baselines on the same 24-site instance (Table 6): (1) Manual Planning based on proximity zones that reflects actual 2024 administrative practice; (2) a Load-

Balanced Manual Plan produced by an experienced coordinator who attempted to balance workload by inspection; (3) Random Cover plus Random Assignment over 30 seeds; (4) Greedy Cover plus Random Assignment over 50 seeds; and (5) a Simple Greedy policy using Greedy Cover plus Longest-Processing-Time (LPT) assignment without the equal-main-trip constraint. All methods were run on the same hardware (Apple M3, 16 GB RAM, Python 3.10). The proposed heuristic matches the best feasible solution in distance (2,399.5 km) and workload variance (0.58), while strictly satisfying both the equal-main-trip constraint and the $\max - \min \leq 3$ policy bound. Random assignments produced expected workload gaps of 32 to 41 points, two orders of magnitude worse than the proposed method, confirming that workload equity cannot emerge from chance. Manual plans, even expert load-balanced ones, produced 8–18% more driving distance and 1–2 additional trips because human planners cannot enumerate all feasible ordered triples.

3.4 Sensitivity Analysis of the Equity Threshold

Table 7 reports the sensitivity of the solution to the equity threshold τ in $\max - \min \leq \tau$, $\tau \in \{1, 2, 3, 4\}$. The branch-and-bound search was run to exhaustion for $\tau = 1$ and with early termination once the target gap was met for $\tau \in \{2, 3, 4\}$. The total driving distance is unaffected because it is fixed in Phase 2 and decoupled from Phase 3.

Table 7. Sensitivity to the equity threshold τ .

Threshold τ	Feasible?	Best Workload Gap	Variance	Total Distance (km)	Runtime (s)
1	No (exhaustive search proved infeasible)	2 (optimum)	0.58	2,399.5	11.89
2	Yes	2	0.58	2,399.5	0.30
3	Yes	2 / 3	0.58 – 0.92	2,399.5	0.28
4	Yes	2 / 3	0.58 – 0.92	2,399.5	0.28

The analysis yields two findings. First, $\tau = 1$ is infeasible on the studied instance: the total points (93) modulo six (15.5) preclude an exactly equal partition, so the theoretical minimum gap is 1 only if trips could be split, which is not allowed; the achievable minimum under the equal-main-trip constraint is 2. Second, thresholds $\tau \geq 2$ are all feasible, and the proposed method reaches the optimum gap of 2 whenever allowed by the search strategy. The institutional policy $\tau \leq 3$ is therefore consistently satisfiable with a comfortable margin of one unit.

3.5 Scalability Analysis

To assess scalability beyond the operational dataset, synthetic instances of $n = 50$ and $n = 100$ placement sites were generated by augmenting the real matrix with randomly drawn symmetric distances in [5, 180] km, visiting times in [46, 210] min, and workload points in $\{1, \dots, 11\}$. Table 8 reports the end-to-end runtime and the number of candidate trips. $|C|$, selected trips $|K|$, driving distance, and LPT-based $\max - \min$ gap. All three instance sizes solve in under one second on a standard workstation.

Table 8. Scalability of the three-phase heuristic.

n (locations)	$ C $ (candidates)	$ K $ (selected trips)	Total Distance (km)	$\max - \min$	Runtime (s)
24	861	12	2,399.5	2	0.014
50	5,046	24	4,738.0	2	0.107
100	22,461	47	9,738.9	4	0.923

Runtime scales near-linearly with $|C|$, which in turn is bounded by $O(n^3)$. The selected-trip count $|K|$ approximately doubles as n doubles, consistent with the 400-minute minimum trip duration. The $\max - \min$

gap grows modestly because the LPT fallback loses the global optimality guarantee at larger $|K|$; incorporating a small branch-and-bound pass for $|K| \leq 14$ clusters of trips remain an effective extension.

3.6 Robustness Analysis

Real driving times are subject to traffic, weather, and minor routing changes that the static matrix D cannot capture. To test robustness, each entry of D was multiplied by a symmetric random factor drawn from $U(1 - \alpha, 1 + \alpha)$ for $\alpha \in \{0, 5\%, 10\%\}$ and the entire heuristic was re-run. For each noise level, 20 Monte Carlo [20] replications were performed; the true (unperturbed) driving distance was then measured on the resulting plan to ensure fair comparisons.

Table 9. Robustness under distance noise (20 replications per level).

Noise Level α	Mean Distance (km) on true D	Std. Dev. (km)	Mean max – min	Mean Trip Count
0%	2,399.5	0.00	2.00	12.00
$\pm 5\%$	2,334.9	142.69	3.60	11.50
$\pm 10\%$	2,333.4	124.22	3.25	11.55

Two observations emerge. First, the mean driving distance on the true matrix D actually decreases slightly under noise, implying that the greedy selection on the unperturbed matrix is not globally optimal and that perturbations occasionally help the heuristic escape local inefficiencies — an encouraging signal that the method is not brittle. Second, the workload gap max – min inflates from 2 to about 3.4 on average, remaining well within the institutional $\tau \leq 4$ bound. The standard deviation of distance (about 5–6% of the mean) matches the input noise magnitude, confirming a well-behaved linear response.

3.7 Greedy vs. Random Set Cover

To justify the use of greedy set cover, we compared it against a random feasible cover strategy over 100 seeds. Table 10 summarizes the results.

Table 10. Greedy vs. random set cover (100 random seeds).

Strategy	Trips $ K $ (mean)	Trips (max)	Distance mean (km)	Distance min (km)
Greedy Set Cover	12	12	2,399.5	2,399.5
Random Set Cover	17.53	21	4,282.8	3,262.9

Greedy covering reduces the trip count. $|K|$ by 31% and the driving distance by 44% relative to the best random cover observed across 100 runs, empirically confirming the theoretical $H(n)$ bound and justifying the design choice.

3.8 Sustainability Analysis

Sustainability is examined through three complementary proxies. Environmentally, the 2,399.5 km driven over 12 trips can be translated into fuel consumption and CO₂ emissions using typical passenger-vehicle coefficients (0.08 L/km for a petrol sedan [21] and 2.31 kg CO₂/L [22]), yielding roughly 191.9 liters of fuel and 443 kg of CO₂ — a 13.8% reduction compared with the 2,781 km baseline manual plan. Economically, at the reference fuel price of 35 THB/L, fuel cost savings exceed 1,800 THB per supervision cycle, which is significant when multiplied across academic years. Socially, sustainability is captured by fairness in workload. The variance of the final workload points W_f across the six supervisors is 0.58, versus 4.17 for a load-balanced manual plan and 6.67 for the proximity-zone manual plan. Reduced variance lowers the risk of burnout among over-supervised faculty and aligns with institutional equity policy, supporting longer-term retention of Co-op staff. Operationally, each supervisor spends a consistent 6–9 hours per supervisory day, with minimal idle time, improving efficiency without compromising assessment quality.

3.9 Discussion

The empirical results converge with and, in three specific respects, extend the existing VRP literature. Consistent with Braekers et al. [5] and the balanced-CVRP literature [18], we confirm that greedy heuristics remain a strong default for small-to-medium operational instances even when metaheuristic and learning-based methods are available. Consistent with Matl et al. [8], we observe that a soft balancing penalty could have produced lower variance at the cost of violating the hard institutional bound — reinforcing their argument that practitioners should choose the objective form that matches operational semantics rather than defaulting to sum-of-squares. The results extend the literature in three ways. First, we provide a concrete academic-supervision formulation in which the workload W_f is discrete (points) rather than continuous (length), and shows that discrete balancing changes the feasibility frontier: $\tau = 1$ is infeasible because of parity, even though the distances would allow a perfectly balanced partition in the continuous case. Second, we introduce an equal-main-trip constraint $\sum_k x_{k,f} = |K|/|F|$ that is rarely enforced in logistics but essential in educational settings; our experiments demonstrate that enforcing it adds only 0.3 s to the total runtime. Third, we report robustness to distance noise at a level rarely tested in VRP studies; the heuristic is shown to be stable at $\pm 10\%$, an unusual but practically important property given the absence of reliable real-time traffic data.

A seemingly contradictory finding is that the Simple Greedy LPT baseline reaches the same max – min gap as the proposed method on this instance. This does not contradict the value of the branch-and-bound assignment: LPT is a well-known 4/3-approximation that is tight in the worst case and provides no guarantee on the equal-main-trip constraint. On larger-workload-variance instances (Section 3.5, $n = 100$), LPT already starts to drift (gap 4 vs. 2 at $n = 24$), and the lack of constraint enforcement would surface in any instance with an indivisible trip count. Thus, LPT is a valid practical fallback but not a replacement for the constrained exact solver.

3.10 Limitations

The study has five acknowledged limitations. (1) The real instance contains only $n = 24$ sites; although it reflects the true operational scale of a single program, generalisability to larger multi-program scenarios has been tested only on synthetic extensions. (2) The method is heuristic; it does not produce an optimality certificate except in the trivial cases handled by branch-and-bound. (3) The driving-distance matrix D is static and does not account for traffic, weather, or time-of-day effects. (4) Discrete points represent workload p_i aggregated from students, mentoring sessions, and assessment depth; other qualitative factors (supervisor expertise, language preferences, mentor–mentee fit) are out of scope. (5) The equity threshold τ is a policy parameter: its value encodes institutional fairness norms and may need recalibration across faculties, departments, or national contexts.

4. Conclusions

This study proposes and validates a three-phase heuristic that jointly solves the Set Partitioning Vehicle Routing Problem with Equitable Workload (SPVRP-EW) for cooperative-education supervision. Applied to a real instance of $n = 24$ placement sites across four Thai provinces, the system produced $|K| = 12$ trips totaling 2,399.5 km with a workload disparity of only 2 points, comfortably within the institutional policy bound of $\tau = 3$. Extensive comparisons against manual, random, and simple-greedy baselines demonstrate superior logistical and equity performance. Sensitivity analysis shows that the heuristic attains the optimum under any practical equity threshold $\tau \geq 2$. Scalability tests up to $n = 100$ synthetic sites confirm sub-second runtime. Robustness tests under ± 5 – 10% distance noise confirm a well-behaved linear response. The system, therefore, constitutes a practical, interpretable decision-support tool for program administrators that simultaneously advances environmental, social, and operational sustainability of Cooperative Education supervision.

5. Future Work

Three extensions are planned. First, integrating real-time traffic and time-dependent VRP formulations [6] to replace the static distance matrix D with live estimates, allowing dynamic re-planning.

Second, scaling the method to large multi-program or multi-campus scenarios by coupling the Phase-2 greedy cover with a Large Neighborhood Search or Genetic Algorithm [23, 24] for candidate refinement, and by replacing the Phase-3 branch-and-bound with a decomposition-based exact solver. Third, embedding the heuristic in an interactive web application and calibrating its equity parameter τ against multi-institution faculty-load policy data to test external validity.

6. Acknowledgments

This research was supported by the Faculty of Informatics, Burapha University, Thailand.

Author Contributions: Conceptualization, A.S. and J.A.; methodology, A.S.; validation, J.A. and A.S.; formal analysis, J.S.; investigation, A.S.; resources, J.S.; data curation, J.S.; writing—original draft preparation, A.S.; writing—review and editing, J.A.; visualization, J.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Faculty of Informatics, Burapha University, Thailand.

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

References

- [1] Jackson, D. Student perceptions of the importance of employability skill provision in business undergraduate programs. *J. Educ. Bus.* **2013**, *88*(5), 271–279. <https://doi.org/10.1080/08832323.2012.697928>
- [2] Billett, S. Learning through work: Exploring instances of relational interdependencies. *Int. J. Educ. Res.* **2008**, *47*(4), 232–240. <https://doi.org/10.1016/j.ijer.2008.07.006>
- [3] Smith, C.; Ferns, S.; Russell, L.; Cretchley, P. *The Impact of Work Integrated Learning on Student Work-Readiness*; Office for Learning and Teaching: Australia, **2014**. <https://hdl.voced.edu.au/10707/337518>
- [4] Harris, M. Academic workloads: Achieving equity and flexibility. *J. Tertiary Educ. Adm.* **1993**, *15*(2), 141–155. <https://doi.org/10.1080/1036970930150202>
- [5] Braekers, K.; Ramaekers, K.; Van Nieuwenhuysse, I. The vehicle routing problem: State of the art classification and review. *Comput. Ind. Eng.* **2016**, *99*, 300–313. <https://doi.org/10.1016/j.cie.2015.12.007>
- [6] Adamo, T.; Gendreau, M.; Ghiani, G.; Guerriero, E. A review of recent advances in time-dependent vehicle routing. *Eur. J. Oper. Res.* **2024**, *319*(1), 1–15. <https://doi.org/10.1016/j.ejor.2024.06.016>
- [7] Pillac, V.; Gendreau, M.; Guéret, C.; Medaglia, A. L. A review of dynamic vehicle routing problems. *Eur. J. Oper. Res.* **2013**, *225*(1), 1–11. <https://doi.org/10.1016/j.ejor.2012.08.015>
- [8] Matl, P.; Hartl, R. F.; Vidal, T. Workload equity in vehicle routing problems: A survey and analysis. *Transp. Sci.* **2018**, *52*(2), 239–260. <https://doi.org/10.1287/trsc.2017.0744>
- [9] Lehuédé, F.; Péton, O.; Tricoire, F. A lexicographic minimax approach to the vehicle routing problem with route balancing. *Eur. J. Oper. Res.* **2020**, *282*(1), 129–147. <https://doi.org/10.1016/j.ejor.2019.09.010>
- [10] Jozefowiez, N.; Semet, F.; Talbi, E.-G. The bi-objective covering tour problem. *Comput. Oper. Res.* **2007**, *34*(7), 1929–1942. <https://doi.org/10.1016/j.cor.2005.07.022>
- [11] Halvorsen-Weare, E. E.; Savelsbergh, M. W. P. The bi-objective mixed capacitated general routing problem with different route balance criteria. *Eur. J. Oper. Res.* **2016**, *251*(2), 451–465. <https://doi.org/10.1016/j.ejor.2015.11.024>
- [12] Kaewploy, W.; Boonkanok, R.; Lertwachara, K. One-day trip itinerary planning for visitors to Songkhla City. *ASEAN J. Sci. Technol. Rep.* **2025**, *28*(6), 33–45. <https://doi.org/10.55164/ajstr.v28i6.259060>
- [13] Chvátal, V. A greedy heuristic for the set-covering problem. *Math. Oper. Res.* **1979**, *4*(3), 233–235. <https://doi.org/10.1287/moor.4.3.233>
- [14] Johnson, D. S. Approximation algorithms for combinatorial problems. *J. Comput. Syst. Sci.* **1974**, *9*(3), 256–278. [https://doi.org/10.1016/S0022-0000\(74\)80044-9](https://doi.org/10.1016/S0022-0000(74)80044-9)
- [15] Korf, R. E. A complete anytime algorithm for number partitioning. *Artif. Intell.* **1998**, *106*(2), 181–203. [https://doi.org/10.1016/S0004-3702\(98\)00086-1](https://doi.org/10.1016/S0004-3702(98)00086-1)

- [16] Della Croce, F.; Scatamacchia, R. The Longest Processing Time rule for identical parallel machines revisited. *J. Sched.* **2020**, 23(2), 163–176. <https://doi.org/10.1007/s10951-018-0597-6>
- [17] Wang, X.; Golden, B.; Wasil, E. The min–max multi-depot vehicle routing problem: Heuristics and computational results. *J. Oper. Res. Soc.* 2015, 66(9), 1430–1441. <https://doi.org/10.1057/jors.2014.108>
- [18] Borgulya, I. An algorithm for the capacitated vehicle routing problem with route balancing. *Cent. Eur. J. Oper. Res.* **2008**, 16(4), 331–343. <https://doi.org/10.1007/s10100-008-0062-2>
- [19] Kool, W.; van Hoof, H.; Welling, M. Attention, learn to solve routing problems! In *Proceedings of the International Conference on Learning Representations (ICLR)*, **2019**. <https://arxiv.org/abs/1803.08475>
- [20] Doubilet, P.; Begg, C. B.; Weinstein, M. C.; Braun, P.; McNeil, B. J. Probabilistic sensitivity analysis using Monte Carlo simulation: A practical approach. *Med. Decis. Making* **1985**, 5(2), 157–177. <https://doi.org/10.1177/0272989X8500500205>
- [21] Department of Alternative Energy Development and Efficiency (DEDE), Ministry of Energy, Thailand. *Thailand Energy Statistics 2022*; DEDE: Bangkok, **2022**.
- [22] Department for Environment, Food & Rural Affairs (DEFRA), UK. *UK Government GHG Conversion Factors for Company Reporting: 2024*; DEFRA: London, **2024**.
- [23] Pisinger, D.; Ropke, S. Large neighborhood search. In *Handbook of Metaheuristics*; Springer: Cham, Switzerland, **2019**; pp 99–127. https://doi.org/10.1007/978-3-319-91086-4_4
- [24] Vidal, T.; Crainic, T. G.; Gendreau, M.; Prins, C. A hybrid genetic algorithm with adaptive diversity management for a large class of vehicle routing problems with time-windows. *Comput. Oper. Res.* **2013**, 40(1), 475–489. <https://doi.org/10.1016/j.cor.2012.07.018>