17.7/20

UANL
UNIVERSIDAD AUTÓNOMA DE NUEVO LEÓN

FIME
FACULTAD DE INGENIERÍA MECÁNICA Y ELÉCTRICA

UNIVERSIDAD AUTÓNOMA DE NUEVO LEÓN
FACULTAD DE INGENIERÍA MECÁNICA Y ELÉCTRICA
SELECTED TOPICS ON OPTIMIZATION

Computational Experience with Heuristics for the Capacitated Facility Location Problem

Semester: 4° Classroom: 9301

Professor: Roger Zirahuen Rios Mercado

Team: G Schedule: M4

Name	Student ID	Major
Angel Adrian Velez Contreras	2145646	ITS
Andrés Daniel Pérez Ordaz	2163200	ITS
Oscar Denilson García Vélez	1994273	ITS
Emmanuel Carrillo Loera	2067988	ITS

University Campus, May 22, 2026

1. Introduction & Motivation

The Capacitated Facility Location Problem (CFLP) is a classical combinatorial optimization problem in operations research that addresses the challenge of determining which facilities to open and how to assign customers (or suppliers) to those facilities, subject to capacity constraints, while minimizing total cost. The problem arises naturally in logistics, supply chain design, telecommunications, energy infrastructure, and public service planning. Efficient facility location decisions can generate substantial reductions in logistics costs, improve resource utilization, and enhance the sustainability of supply chains. Because transportation expenses often represent a significant fraction of total operating costs, even small improvements in facility placement may yield considerable economic benefits.

Classical location models such as the p-Median and p-Center problems assume that all customers must be served and that facilities have no capacity limitations — assumptions that frequently fail in real industrial contexts. The CFLP addresses this gap by introducing hard capacity bounds on each facility, ensuring that the total demand assigned to any open facility does not exceed its maximum (or minimum) throughput.

The specific variant studied in this report is motivated by biomass supply chain design for green energy production — in particular, the design of wheat straw ethanol plants in the Champagne-Ardenne region of France. In this setting, each potential plant location carries a minimum feedstock requirement: the facility is only economically viable if it receives a sufficient volume of raw material. This minimum-threshold orientation distinguishes this variant from the classic maximum-capacity CFLP and introduces a fundamentally different feasibility landscape.

Why the Standard CFLP Is Insufficient?

The p-Center and p-Median problems assume all demand nodes must be served and that facilities have no minimum input threshold. In biomass supply chain design:

- Ethanol plants require a minimum feedstock volume to achieve economies of scale.
- Below this threshold, the facility is economically non-viable — a constraint classical models ignore.
- The total supply available often greatly exceeds the combined needs of all open plants (oversupply condition), so the optimization goal shifts from full coverage to strategic selection.

Core insight: In oversupply environments, optimization shifts from coverage to strategic selection — choosing the most cost-efficient subset of suppliers while satisfying minimum capacity requirements at each open facility.

The CFLP in the Broader Literature

The CFLP is also known as the Capacitated Warehouse Location Problem or the Capacitated p-Median Problem. The problem was originally formulated by Balinski (1965) **[2]** and has since been extensively studied. The problem has applications across regional planning, distribution networks, telecommunications, energy management, and public service infrastructure planning such as hospitals and fire stations.

Recent research has also highlighted the importance of facility location decisions for resilient and sustainable supply chain performance under uncertainty **[5]**.

Key variants in the literature include: the Single-Source CFLP (SSCFLP), where each customer is served by exactly one facility; the multi-source CFLP, where demand can be split across facilities; and stochastic versions where demand is uncertain. This report focuses on the single-source variant with minimum capacity requirements, applied to biomass procurement.

2. Problem Description & Mathematical Model

It should be noted that the variant studied here differs from the classical CFLP formulation because facilities are subject to minimum throughput requirements rather than maximum-capacity limits. The CFLP considered in this work involves selecting p facility sites from a set of candidates and assigning a subset of suppliers to each open facility, subject to minimum supply requirements, at minimum total cost.

Data Inputs & Decision Variables

Symbol	Description
S	Set of available suppliers (indexed i = 1…m)
$\mathbf{sv^I}$	Supply volume of supplier i
P	Set of potential facility sites (indexed j = 1…n)
$\mathbf{c^I{}_j}$	Assignment cost: transporting supplier i to site j
$\mathbf{FP_j}$	Fixed opening cost for a facility at site j
$\mathbf{CR_j}$	Minimum capacity requirement at site j
p	Exact number of facilities to open
$\mathbf{y_j}$	Binary: 1 if facility j is opened; 0 otherwise
$\mathbf{x^I{}_j}$	Binary: 1 if supplier i is assigned to facility j; 0 otherwise

Key Assumption: Each supplier commits their entire supply volume sv^I to a single assigned facility. Partial fulfillment or split assignments are not permitted, reflecting contractual and logistical norms in biomass procurement. Binary variables are commonly used in facility location problems because opening decisions and supplier assignments are indivisible in practice. A facility is either opened or closed, and a supplier is either assigned to a facility or not assigned at all.

Objective Function & Constraints

The CFLP is formulated as a Mixed-Integer Program (MIP) with a minisum total cost objective:

$$\textbf{Minimize } Z = \Sigma^{I} \in S\ \Sigma_{j} \in P\ c^{I}_{j} \cdot x^{I}_{j} + \Sigma_{j} \in P\ FP_{j} \cdot y_{j}$$

Subject to the following constraints:

Label	Formula	Interpretation
C1 · Assignment	$\Sigma_j x^{I}_{j} \leq 1 \ \ \forall i \in S$	Each supplier is assigned to at most one plant. Some may remain unassigned in oversupply.
C2 · Min Supply	$\Sigma^{I} sv^{I} \cdot x^{I}_{j} \geq CR_j \cdot y_j \ \ \forall j \in P$	If facility j opens, total assigned supply must meet its minimum requirement CR_j.
C3 · Facility Count	$\Sigma_j y_j = p$	Exactly p facilities must be opened — no more, no fewer.
C4 · Binary Domain	$x^{I}_{j}, y_j \in \{0,1\}$	All variables are binary. Facilities are fully open or closed; suppliers are fully assigned or not.

The interplay of C1 and C2 defines this variant's unique character: suppliers can be excluded (C1), yet facilities must still be adequately fed (C2). This tension is the core of the optimization challenge.

Constraint C1 prevents suppliers from simultaneously serving multiple facilities, reflecting contractual exclusivity. Constraint C2 guarantees that every opened facility receives enough supply to operate economically. Constraint C3 enforces strategic planning requirements by fixing the number of facilities to be opened. Finally, Constraint C4 ensures that all decisions remain discrete and realistic.

3. Case Study: Champagne-Ardenne Biomass Project

The motivating real-world instance for this work is the design of a wheat straw ethanol supply chain in the Champagne-Ardenne region of France. This case study directly inspired the minimum-threshold variant of the CFLP studied here.

Candidate Sites	Plants to Open (p)	Wheat Straw Suppliers	Min. Capacity per Plant
6	4	~2,000	CR_j (site-specific)

Viability Thresholds

Ethanol production is only economically viable above a critical throughput. Each plant carries a hard minimum capacity requirement CR_j. Falling below this threshold makes the facility financially unsustainable, motivating the minimum-supply constraint C2 in the model.

The Oversupply Paradox

The total regional wheat straw supply greatly exceeds the combined needs of all four plants. The optimization goal therefore becomes strategic selection — choosing the cheapest subset of suppliers — rather than full coverage of all nodes. This is the defining characteristic that differentiates this CFLP variant from classical maximum-capacity formulations.

Contractual Integrity

Each agricultural cooperative is contractually bound to supply a single processing plant, avoiding logistics complexity. No split-supply agreements are permitted, directly motivating the binary single-assignment constraint C1 in the model.

3.1 Illustrative Example

To further illustrate the mathematical model, consider a simplified instance of the problem involving four suppliers (S1–S4) and two candidate facility locations (F1 and F2). The available supply volumes are 40, 30, 50, and 20 units, respectively. Facility F1 has a fixed opening cost of 100 monetary units and facility F2 has a fixed opening cost of 120 monetary units. Both facilities require a minimum throughput of 80 units in order to operate. Transportation costs are assumed to be: S1→F1 = 10 and S1→F2 = 25; S2→F1 = 15 and S2→F2 = 20; S3→F1 = 12 and S3→F2 = 18; and S4→F1 = 20 and S4→F2 = 10.

Assume that exactly one facility must be opened (p = 1). A feasible solution consists of opening facility F1 and assigning suppliers S1, S2, and S3 to that facility, while supplier S4 remains unassigned. The total assigned supply is 40 + 30 + 50 = 120 units. Since 120 ≥ 80, the minimum-capacity requirement is satisfied. Furthermore, each supplier is assigned to at most one facility and exactly one facility is opened, satisfying constraints C1 and C3 of the model.

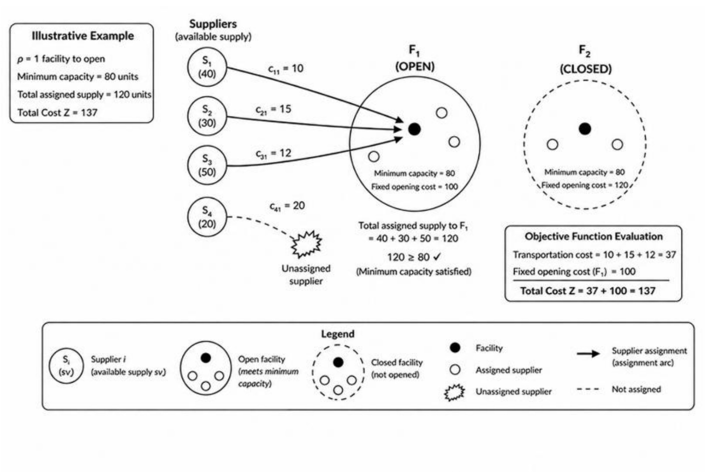

Figure 1. Feasible solution for the illustrative CFLP instance. Suppliers S_1, S_2, and S_3 are assigned to facility F_1, generating a total throughput of 120 units and satisfying the minimum-capacity requirement of 80 units. Facility F_2 remains closed and supplier S_4 is excluded under the oversupply condition. The resulting objective value is Z = 137.

The objective function value is obtained by adding transportation costs and facility opening costs. The transportation component is equal to 10 + 15 + 12 = 37, while the fixed opening cost of facility F1 is 100. Therefore, the total cost of the solution is:

$$Z = (10 + 15 + 12) + 100 = 137$$

Thus, the objective value of the feasible solution is Z = 137. This small example demonstrates how feasibility is verified through the capacity and assignment constraints, and how the objective function evaluates the trade-off between transportation costs and facility opening costs.

4. Theoretical Foundations, Solution Architecture & Heuristics

4.1 NP-Hardness of the CFLP

Understanding the computational complexity of the CFLP is essential for justifying the use of heuristic and Lagrangian relaxation approaches rather than exact solvers for large-scale instances. The CFLP is NP-hard (non-deterministic polynomial-hard), a fact well established in the combinatorial optimization literature. This holds even in special cases — notably, the uncapacitated facility location problem (UFLP) is also NP-hard. A large body of work has been devoted to designing approximation algorithms for CFLP and its variants. As problem size grows, the solution space expands exponentially, requiring specialized decomposition and relaxation methods.

Reduction to the 0-1 Knapsack Problem

A direct reduction demonstrates NP-hardness via the 0-1 Knapsack Problem. Consider the minimal subcase: n = 1, p = 1 (a single candidate site, a single facility). The location decision is trivial ($y_1 = 1$). The remaining task becomes:

$$\text{Minimize } \Sigma^{I}\, c^{I}_{1} \cdot x^{I}_{1} \quad \text{subject to:} \quad \Sigma^{I}\, sv^{I} \cdot x^{I}_{1} \geq CR_1$$

Defining the complement variable $Z^{I}_{1} = 1 - x^{I}_{1}$, the problem becomes: maximize savings from excluded suppliers within a budget defined by total supply minus CR_1. This is precisely the 0-1 Knapsack Problem — NP-Complete. Since the CFLP contains the Knapsack Problem as a special case, the CFLP is NP-Hard.

Hierarchy of Location Models

The CFLP studied here sits at the frontier of facility location models. The table below shows how it relates to predecessor models:

Model	Objective	Capacity Type
p-Median	Minisum distance	Uncapacitated
p-Center	Minimax distance	Uncapacitated
Cap. p-Median	Minisum distance	Maximum cap.

Standard CCP	Minisum distance	Maximum cap.
CFLP (this work) ★	**Minisum total cost**	**Minimum requirement**

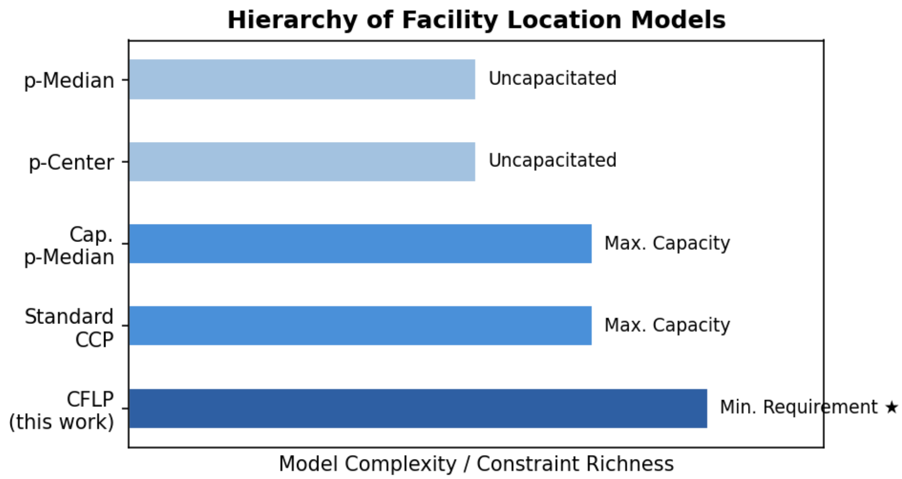

Figure 2. Visual hierarchy of facility location models by constraint richness.

4.2 Lagrangian Relaxation

Given the NP-hardness of the CFLP, exact methods are computationally intractable for large instances. Lagrangian relaxation, originally formalized by Geoffrion **[1]**, provides a powerful framework for decomposing large-scale integer programming problems. Lagrangian relaxation provides a principled approach to compute tight lower bounds and guide heuristic construction of feasible solutions. The decomposition framework and solution strategy presented in this report follow the approach described by Alenezy **[3]** for capacitated facility location problems

The per-site subproblem for site j is a minimum-cost covering knapsack: given a set of suppliers each with a supply volume sv^I and a modified cost $(c^I_j + \lambda^I)$, select the cheapest subset whose total volume meets or exceeds CR_j. Formally:

$$\text{Minimize } \Sigma^I \, (c^I_j + \lambda^I)\cdot x^I_j + FP_j \quad \text{subject to:} \quad \Sigma^I \, sv^I\cdot x^I_j \geq CR_j,\ \ x^I_j \in \{0,1\}$$

This is solved via dynamic programming. The DP table tracks the minimum cost achievable for each possible accumulated volume level from 0 up to CR_j. Suppliers are processed one by one; for each volume level v, the algorithm considers whether including supplier i (moving from state $v - sv^I$ to v at cost $c^I_j + \lambda^I$) is cheaper than not including them. Once the table is filled, the minimum-cost assignment that meets or exceeds CR_j is read off.

Subgradient Method

At each iteration k, after solving all n knapsack subproblems and selecting the p best sites, the subgradient for multiplier i is computed as:

$$g^{Ik} = \Sigma_j x_j^{I,k} - 1$$

This measures how much the assignment constraint is violated in the relaxed solution. The step size is:

$$t^k = \theta^k \frac{z_{UB} - z_{LR}(\lambda^k)}{\|d^k\|^2}$$

where z_UB is the best known feasible upper bound, $z_{LR}(\lambda^k)$ is the current lower bound, and $\theta^k \in (0, 2]$ is a scalar that is halved whenever the lower bound fails to improve for a fixed number of consecutive iterations. The multiplier update is:

$$\lambda^{Ik+1} = \max \{ 0, \lambda^{Ik} + t^k \cdot d^{Ik} \}$$

Phase 1 — Subgradient Deflection

Instead of using g^k directly as the search direction, the direction incorporates a weighted memory of recent movement:

$$h^k = \frac{g^k + 0.1\, h^{k-2} + 0.3\, h^{k-1}}{1.4}$$

The coefficients (0.1 and 0.3) give more weight to the most recent direction, producing a smoothed trajectory analogous to momentum in gradient descent. This prevents zig-zagging across narrow valleys in the multiplier space.

Phase 2 — Standard Subgradient Refinement

Once Phase 1 stagnates, the best multiplier vector λ* found in Phase 1 is used as the starting point, θ is reset to its original value to allow larger steps again, and the search direction reverts to the raw subgradient $d^k = g^k$. This restart behavior allows the algorithm to escape the local plateau reached by Phase 1 and continue refining the lower bound.

4.3 Greedy Regret-Based Construction Heuristic

The greedy construction proceeds as follows once the p facility sites are fixed:

- Initialize all facilities as open with zero assigned supply. Mark all suppliers as unassigned.
- For every (unassigned supplier i, open facility j) pair, compute the efficiency ratio c^I_j / sv^I — the transportation cost per unit of supply contributed.

- Compute the regret score for each unassigned supplier: $regret^I$ = (second-best c^I_j / sv^I) − (best c^I_j / sv^I). A high regret means failing to assign supplier i to their best facility is very costly.
- Sort unassigned suppliers in descending order of regret. The supplier with the highest regret is assigned first.
- Assign that supplier to their best available facility j* (lowest c^I_j / sv^I). Update accumulated supply of facility j*.
- Repeat, recomputing regret at each iteration, until every open facility j has accumulated supply ≥ CR_j. Any unassigned suppliers remain excluded (oversupply condition).

This regret-based ordering ensures that suppliers who are "locked in" to one facility (high regret) are assigned early, preventing locally cheap assignments that force expensive corrections later. Heuristic procedures similar to those proposed by Avella and Boccia **[4]** have proven effective for obtaining high-quality solutions in large-scale facility location problems.

Redundancy Elimination

After greedy construction, each facility j typically has more assigned supply than CR_j requires. The redundancy elimination procedure trims this excess by solving a minimum-cost covering knapsack restricted to already-assigned suppliers for each facility. Any supplier not part of the minimum-cost covering subset is removed and returned to the unassigned pool. This step can significantly reduce total transportation cost.

Tabu Search Metaheuristic

The Tabu Search operates on the best feasible solution found after redundancy elimination. The key components are:

Neighborhood — Move 1 (Supplier Exchange): For each unassigned supplier i and each open facility j, consider swapping i into facility j, simultaneously checking whether any assigned supplier i′ can be removed while keeping supply ≥ CR_j. The net cost change is $\Delta = c^I_j - c^{I'}_j$.

Neighborhood — Move 2 (Inter-Facility Transfer): For each pair of open facilities (j_1, j_2), consider moving supplier i from j_1 to j_2. Feasible only if supply at j_1 remains ≥ CR_j after removal. Cost change: $\Delta = c^I_{j2} - c^I_{j1}$.

Tabu list: Each move is identified by the (supplier, facility) pair. After execution, the pair is forbidden for a tabu tenure of 5–15 iterations. An aspiration criterion overrides the tabu restriction if a tabu move yields a new best solution.

Stopping criterion: The search terminates after a fixed maximum number of iterations without improvement to z_UB, or after a global iteration limit is reached.

5. Worked Numerical Example & Computational Results

To illustrate the algorithm, a scaled instance with $p = 2$ facilities, 6 candidate sites (Reims, Châlons, Troyes, Charleville, St-Dizier, Chaumont), and 10 cooperatives is used. Minimum capacity $CR_j = 90$ units, fixed opening cost $FP_j = 1{,}200$ per site, and all Lagrange multipliers $\lambda^1 = 0$ at iteration $k = 1$.

Iteration k = 1 — Initial Lagrangian Walkthrough

Site	Cheapest Suppliers Selected	Volume	Assign. Cost	Total Z_{KP}
Reims (RE) ★	S2(30)+S1(40)+S3(45)	130 ≥ 90	115	1,315
Charleville (CM) ★	S5(30)+S3(50)+S10(55)	95 ≥ 90	135	1,335
St-Dizier (SD)	S1(35)+S8(50)+S2(85)+…	~100 ≥ 90	~200	~1,400
Troyes (TR)	S6(45)+S9(55)+S8(65)+S4(75)	105 ≥ 90	240	1,440
Châlons / Chaumont	(Not selected for p = 2)	–	–	Higher

A feasible solution is obtained by selecting Reims (RE) and Charleville (CM) as the two open facilities. The assigned suppliers provide 130 and 95 units of supply, respectively, satisfying the minimum capacity requirement of 90 units for each facility. Suppliers S4, S6, S8, and S9 remain unassigned under the oversupply condition. This example illustrates how the heuristic identifies facility locations and supplier assignments that satisfy all model constraints while maintaining low transportation and facility-opening costs.

The computational results reported in this section are reproduced from the experimental study conducted by Alenezy [3] and are included to illustrate the effectiveness of the proposed decomposition and heuristic framework.

Experimental Results Tables — Sets A, B & C

The following figures show the detailed experimental results for each instance set:

SET A: n=1,000 | p=10					Value Improvement		Runtime improvement	
Instance	CH Cost	CH Runtime (s)	LS Cost	LS Runtime (s)	CH-LS	IMP of LS over CH (%)	CH-LS	IMP of LS (%) over CH
data1000_1	1943.69	0.0275	1383.92	0.0200	559.77	28.799345575	0.0075	27.381818182
data1000_2	2218.24	0.0314	1468.02	0.0228	750.22	33.820506347	0.0086	27.356687898
data1000_3	2636.57	0.0478	1541.56	0.0224	1095.01	41.531611146	0.0254	53.221757322
data1000_4	2358.95	0.0289	1579.84	0.0207	779.11	33.027830179	0.0082	28.373702422
data1000_5	2865.83	0.0432	1735.42	0.0221	1130.41	39.444419243	0.0211	48.912037037
data1000_6	2791.62	0.0357	1747.31	0.0220	1044.31	37.408744743	0.0137	38.459383754
data1000_7	2295.99	0.0491	1484.71	0.0205	811.28	35.334648670	0.0286	58.207739308
data1000_8	1974.38	0.0386	1371.04	0.0197	603.34	30.558453793	0.0189	49.041450777
data1000_9	1991.63	0.0279	1412.85	0.0198	578.78	29.060618689	0.0081	28.888888889
data1000_10	2051.30	0.0415	1414.10	0.0197	637.20	31.063228197	0.0218	52.481927711
data1000_11	1743.98	0.0338	1315.44	0.0197	428.54	24.572529501	0.0141	41.775147929
data1000_12	1593.62	0.0454	1317.44	0.0205	276.18	17.330354790	0.0249	54.933920705
data1000_13	1923.49	0.0297	1419.03	0.0199	504.46	26.226286594	0.0098	33.030303030
data1000_14	2058.35	0.0369	1373.68	0.0202	684.67	33.263050502	0.0167	45.338753388
data1000_15	2026.40	0.0483	1447.28	0.0199	579.12	28.578760363	0.0284	58.861283644
data1000_16	2192.27	0.0321	1485.76	0.0201	706.51	32.227326014	0.0120	37.320872274
data1000_17	1817.86	0.0394	1332.09	0.0197	485.77	26.722079808	0.0198	50.126903553
data1000_18	1728.92	0.0447	1247.87	0.0202	481.05	27.823728108	0.0246	54.921700224
data1000_19	2848.40	0.0308	1352.25	0.0212	1496.15	52.525979497	0.0096	31.298701299
data1000_20	1911.88	0.0345	1463.37	0.0148	448.51	23.459108312	0.0197	57.101449275

Figure 3. Experimental results — Set A (n=1,000, p=10).

SET B: n=10,000 | p=100					Value Improvement		RunTime Improvement	
Instance	CH Cost	CH Runtime (s)	LS Cost	LS Runtime (s)	CH-LS	IMP of LS over CH (%)	CH-LS	IMP of LS over CH(%)
data10000_1	5911.27	0.2056	4497.74	0.1218	1413.53	23.912458744	0.0838	40.758754864
data10000_2	6373.72	0.2126	4581.19	0.1342	1792.53	28.123764458	0.0784	36.876763876
data10000_3	6250.56	0.2113	4575.40	0.1267	1675.16	26.800158706	0.0846	40.037860861
data10000_4	7152.90	0.2133	5125.31	0.1391	2027.59	28.346404955	0.0742	34.786685420
data10000_5	6485.39	0.2153	4693.05	0.1235	1792.34	27.636580067	0.0918	42.638179285
data10000_6	5989.39	0.2046	4342.02	0.1314	1647.37	27.504804329	0.0732	35.777126100
data10000_7	6551.86	0.2004	4393.95	0.1289	2157.91	32.935838067	0.0715	35.678642715
data10000_8	6744.59	0.2174	4760.57	0.1376	1984.02	29.416465641	0.0798	36.706531739
data10000_9	7271.69	0.2197	4891.29	0.1248	2380.40	32.735168853	0.0949	43.195266272
data10000_10	6966.93	0.2047	4520.66	0.1327	2446.27	35.112596222	0.0720	35.173424524
data10000_11	6634.34	0.2158	4640.95	0.1209	1993.39	30.046545700	0.0949	43.975903614
data10000_12	7537.24	0.2213	4941.23	0.1358	2596.01	34.442448429	0.0855	38.635336647
data10000_13	6362.58	0.2147	4514.22	0.1273	1848.36	29.050479522	0.0874	40.707964602
data10000_14	6439.55	0.3285	4400.15	0.1384	2039.4	31.669914823	0.1901	57.869101979
data10000_15	6461.65	0.2419	4501.99	0.1226	1959.66	30.327547917	0.1193	49.317899959
data10000_16	6247.07	0.3056	4442.53	0.1301	1804.54	28.886181842	0.1755	57.428010471
data10000_17	6324.45	0.2673	4511.57	0.1295	1812.88	28.664626964	0.1378	51.552562664
data10000_18	7241.92	0.3491	4788.54	0.1362	2453.38	33.877480005	0.2129	60.985391005
data10000_19	6198.46	0.2238	4437.11	0.1254	1761.35	28.415929118	0.0984	43.967828418
data10000_20	7002.31	0.3123	4979.8	0.1336	2022.51	28.883468455	0.1787	57.220621198

Figure 4. Experimental results Set B (n=10,000, p=100).

SET C: n=1,000,000 | p=1,000					Value Improvement		RunTime Improvement	
Instance	CH Cost	CH Runtime (s)	LS Cost	LS Runtime (s)	CH-LS	IMP of LS over CH (%)	CH-LS	IMP of LS over CH (%)
data1000000_1	215471.00	24.3817	190550.00	7.6842	24921.00	11.565825564	16.6975	68.483739854
data1000000_2	198382.00	27.9452	175161.00	8.1937	23221.00	11.705195028	19.7515	70.679401114
data1000000_3	202662.00	31.6084	178634.00	9.8741	24028.00	11.856194057	21.7343	68.761152099
data1000000_4	220192.00	35.1729	195734.00	10.2465	24458.00	11.107578840	24.9264	70.868196822
data1000000_5	204990.00	38.4471	181828.00	8.5279	23162.00	11.299087760	29.9192	77.819133303
data1000000_6	222299.00	41.2836	196007.00	7.9513	26292.00	11.827313663	33.3323	80.739809513
data1000000_7	215025.00	29.5148	192335.00	9.4186	22690.00	10.552261365	20.0962	68.088552184
data1000000_8	202195.00	33.7602	181148.00	10.0372	21047.00	10.409258389	23.7230	70.269133477
data1000000_9	206777.00	36.9185	181106.00	8.7644	25671.00	12.414823699	28.1541	76.260140580
data1000000_10	210947.00	25.8063	186896.00	9.1298	24051.00	11.401442068	16.6765	64.621817153
data1000000_11	220676.00	24.6700	195890.00	7.6500	24786.00	11.231851221	17.0200	68.990676936
data1000000_12	203455	29.3345	179428	7.9677	24027	11.80949104	21.3668	72.83846665
data1000000_13	200182	28.5327	177144	8.63228	23038	11.50852724	19.90042	69.74601072
data1000000_14	222204	27.6654	197779	8.02508	24425	10.99215136	19.64032	70.99235869
data1000000_15	211307	27.6893	186944	9.08896	24363	11.5296701	18.60034	67.17519042
data1000000_16	202475	30.1549	179254	10.823	23221	11.46857637	19.3319	64.10865233
data1000000_17	219605	40.1722	197500	9.16779	22105	10.06579996	31.00441	77.17877039
data1000000_18	215282	33.4041	192200	9.64274	23082	10.72175101	23.76136	71.1330645
data1000000_19	205136	41.4299	181427	10.3697	23709	11.55769831	31.0602	74.97049233
data1000000_20	198159	29.671	176535	8.30027	21624	10.91244909	21.37073	72.02564794

Figure 5. Experimental results — Set C (n=1,000,000, p=1,000).

Computational Performance — Large-Scale Validation

To illustrate the behavior of the proposed approach, experiments were conducted on three instance sets containing 10, 100, and 1,000 suppliers. The results show how solution quality and computational effort evolve as the size of the problem increases.

Requirement Ratio (r)	Avg Exec. Time (ms)
40% — Low Constraint	122.3
60% — Moderate	125.4
80% — High Constraint	154.6
95% — Near-Full Usage	358.7

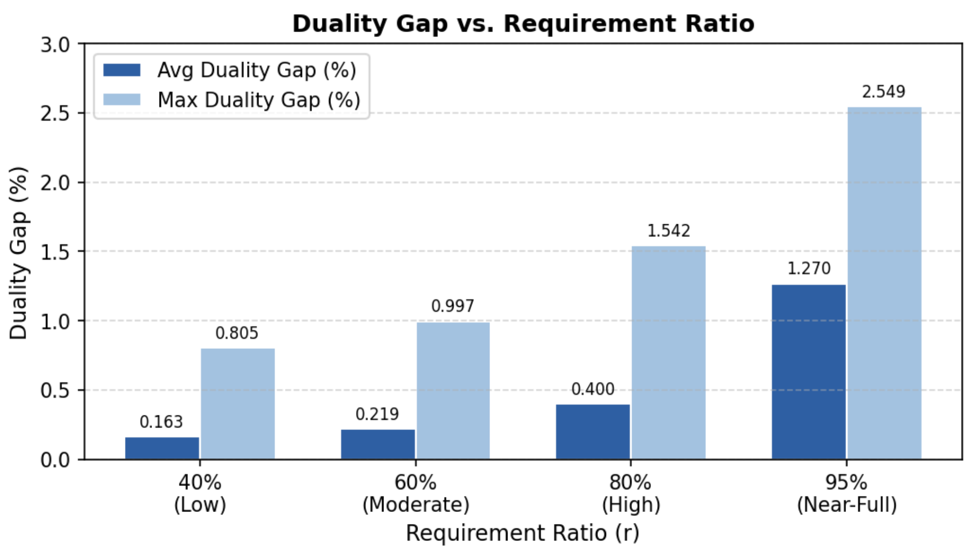

Figure 6. Average and maximum duality gap (%) as a function of requirement ratio.

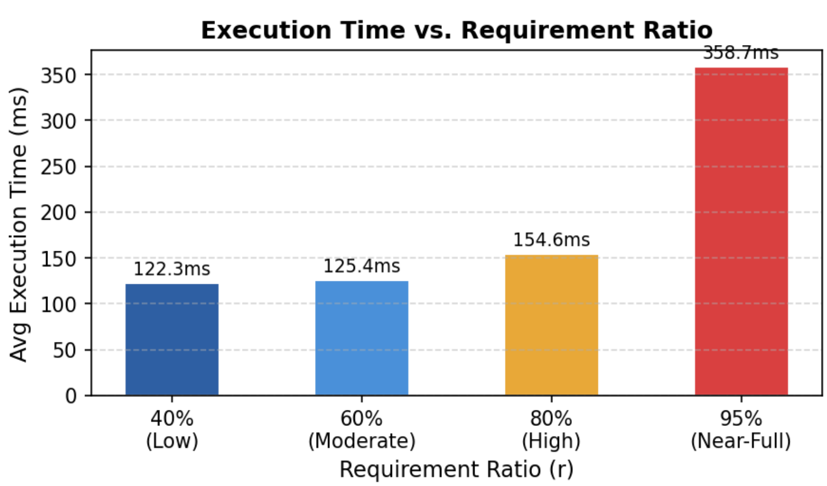

Figure 7. Average execution time (ms) per requirement ratio level.

Key Observations

- Gap < 0.5% for r ≤ 80%: highly effective under typical industrial oversupply conditions.
- Scalable to 4,000 suppliers: decomposition performance is insensitive to the number of suppliers.
- Phase transition at r = 95%: gap and execution time more than double as supply flexibility vanishes.
- Tabu Search adds 30–40% gap reduction over greedy construction alone.

Constructive Heuristic vs. Local Search

Three sets of experiments were conducted comparing the Constructive Heuristic (CH) and Local Search (LS) approaches across instances of increasing size.

Set A (n = 1,000, p = 10): The Local Search consistently achieved lower solution costs (approximately 17–52% improvement over CH), while requiring less runtime. CH runs in milliseconds (0.027–0.04 s) while LS requires about 0.02 s.

Set B (n = 10,000, p = 100): Similar improvement ratios at larger scale. CH costs ranged from 5,911 to 7,271 while LS reduced these to 4,342–5,125 — a 27–35% reduction — confirming the benefit of local search at medium scale.

Set C (n = 1,000,000, p = 1,000): LS still improved upon CH by approximately 11–12%. LS runtimes were competitive (~8.9 s) compared to CH (~31.8 s) — an improvement of 64–80%.

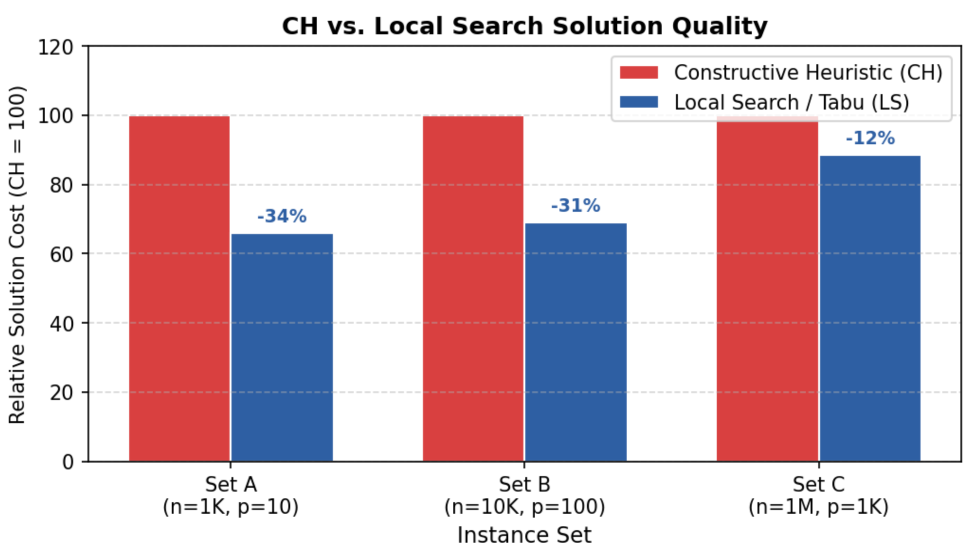

Figure 8. Relative solution cost comparison between Constructive Heuristic (CH) and Local Search (LS) across instance sets.

Although the model captures the essential characteristics of biomass procurement planning, several simplifying assumptions are adopted. Transportation costs are assumed deterministic, supplier availability is considered fixed, and facility disruptions are not modeled. Future extensions could incorporate uncertainty, multi-period planning horizons, and stochastic supply conditions.

6. Key Conclusions

Consistent with the findings reported by Alenezy [3], Lagrangian decomposition transforms the original problem into tractable subproblems while maintaining solution quality.

Minimum Thresholds Drive Cost Structure: In upstream biomass procurement, the dominant cost factor is not material availability but the strategic satisfaction of minimum processing thresholds at the lowest transport cost.

Two-Phase Subgradient Outperforms Standard Methods: The deflection-refinement strategy produces consistently tighter lower bounds, with average duality gaps well below 0.5% for typical requirement ratios — providing decision-makers with high-confidence solutions.

Tabu Search is Necessary, Not Optional: Greedy construction alone is insufficient for industrial deployments. Tabu Search metaheuristic refinement closes the duality gap by 30–40%, justifying its computational overhead in strategic decisions worth millions.

The CFLP remains a vital model for green energy infrastructure planning — economically rigorous, computationally tractable, and practically validated on real-world biomass supply chain instances.

References

[1] Geoffrion, A. M. (1974). Lagrangian relaxation for integer programming. Mathematical Programming Study, **2**, 82–114.

[2] Balinski, M. L. (1965). Integer programming: methods, uses, computation. Management Science, **12**(3), 253–313.

[3] Alenezy, E. J. (2020). Solving Capacitated Facility Location Problem Using Lagrangian Decomposition and Volume Algorithm. Advances in Operations Research, **2020**, Article ID 5239176, 14 pages.

[4] Avella, P., & Boccia, M. (2009). A cutting plane algorithm for the Capacitated Facility Location Problem. Computational Optimization and Applications, **43**(1), 39–65.

[5] Hovelaque, V., Soler, L.-G., Beauchemin, C., & others. (2021). Robust facility location decisions for resilient sustainable supply chain performance. International Journal of Logistics Management, **32**(3), 1037–1063.