Linear Program for EV Charging Station Siting with K-means ℓ_1 Activation Penalty

Abdoulaye Diop
Texas A&M University
College Station, TX, USA
ablaye2002@tamu.edu

Wonhyeok Jang
Texas A&M University
College Station, TX, USA
wjang777@tamu.edu

Thomas Overbye
Texas A&M University
College Station, TX, USA
overbye@tamu.edu

Jonathan Snodgrass
Texas A&M University
College Station, TX, USA
snodgrass@tamu.edu

***Abstract*—This paper presents a feeder-constrained linear programming framework for screening electric vehicle charging station locations on radial distribution feeders. The proposed model selects candidate charging sites while enforcing Kirchhoff's Current Law and branch capacity limits, allowing planners to identify locations that are electrically feasible under feeder-level hosting constraints. To improve spatial relevance, the formulation introduces a continuous activation penalty based on proximity to multi-family housing demand points and k-means demand centroids. This penalty encourages compact siting patterns near concentrated multi-family housing demand without converting the problem into a mixed-integer or nonlinear program. The model is implemented in JuMP with the HiGHS solver and evaluated on four processed feeder-capacity maps derived from PG&E Distribution Capacity / GRIP data. Results show that the method distinguishes capacity-limited feeders from feeders with available headroom, reduces the distance between selected charging sites and multi-family housing demand, and identifies highly utilized feeder segments that may require upgrade investigation. The proposed approach is intended as a fast planning-screening layer rather than a replacement for detailed AC power-flow validation.**

***Index Terms*—EV charging infrastructure, demand location, distribution feeders, linear programming, k-means clustering**

I. Introduction

Public electric-vehicle charging stations (EVCS) must scale rapidly while avoiding overloads on distribution feeders and remaining transparent for planners. Today's siting practice draws from several methodological streams with different objectives and computational trade-offs [1], [2]. National diagnostics of charging coverage [3], demand-forecasting studies [4], and resilience analyses [5] provide valuable inputs but are not siting algorithms.

Prior approaches and limitations. Flow-based siting via the *flow-refueling location model (FRLM)* maximizes driver path coverage [6]–[8] but abstracts the electrical feeder, so grid feasibility is checked post-hoc. GIS/MCDM workflows and metaheuristics such as MOPSO [9] flexibly encode siting criteria but depend on user-chosen weights that affect reproducibility [2]. Feeder-physics formulations (MISOCP, MINLP, two-stage robust) [10]–[12] improve realism but introduce nonconvexity or combinatorial scaling that impedes multi-feeder screening. Recent MILP models [16] retain integer complexity; clustering (k-means) proposes compact candidate sets [13]–[15] but omits in-solve *Kirchhoff's Current Law (KCL)* and per-edge *load hosting capacity (LHC)* constraints. Across these streams, common pain points are: enforcing grid feasibility *in-solve*, sensitivity to weights or seeds, nonconvexity and solver tuning, and data heterogeneity [18].

Research gap. No existing convex screening model simultaneously enforces in-solve KCL, per-edge branch capacity limits, and MFH-biased spatial penalties on radial feeders within a single linear program. This paper addresses that gap. We develop a feeder-aware LP that stays convex while explicitly enforcing radial feeder flow conservation and branch-capacity limits, quickly identifies electrically plausible MFH-targeted siting patterns across many feeders, and flags corridors that repeatedly emerge as bottlenecks—all without mixed-integer or conic complexity. Public synthetic feeder sets such as SMART-DS [17] are reserved as a future AC-validation path.

Contributions. (i) A feeder-constrained EV siting LP enforcing KCL and branch thermal or hosting-capacity limits on a radial network. (ii) A continuous distance-based activation penalty biasing siting toward MFH demand concentrations without altering the model class. (iii) Bottleneck indicators: the LP solve that identifies candidate sites simultaneously flags feeder segments approaching capacity limits for upgrade investigation.

II. Methodology

A. Overview

The workflow runs five stages in sequence: (i) feeder geometry is converted into a directed radial graph rooted at the substation; (ii) MFH demand points are processed and summarized with k-means clustering; (iii) each node receives a distance-based penalty reflecting proximity to MFH buildings and cluster centroids; (iv) MFH demand is attached to nodes via a segment-projection and flow-informed attachment rule; and (v) the unified LP is solved in baseline and penalized modes.

B. Data sources and parameter settings

The study uses four radial feeders from cleaned, processed feeder-capacity maps prepared from PG&E Distribution Capacity / GRIP GIS source layers [19], [20], spanning 269–372 network nodes and 16–82 MFH candidate locations. Per-edge capacities serve as first-order thermal or hosting-capacity screening limits rather than substitutes for full AC

TABLE I
PRINCIPAL MODEL PARAMETERS AND FEEDER 5 STUDY SETTINGS.

Param.	Description	Units	Setting
d_m	MFH EV demand proxy	kW	Level-2 screened demand
d_i	Node aggregate demand	kW	closest MFH demands
C_e	Edge thermal/LHC limit	kW	feeder-specific
λ_0	Base penalty weight	–	7.2
α	Penalty multiplier	–	2, 5, 10, 18, 25
β	Demand relaxation	–	effective max 0.2
$D_{\max}$	MFH attach radius	m	inactive in runs
$\tau_s^\dagger$	Site cutoff	–	$z_i > 0.001$
K	k-means centroids	–	{1,3,5,7} or {1–10}
η	MFH/centroid blend	–	0.6

† Post-processing threshold; not part of LP.

validation [18]; SMART-DS and comparable models are reserved for future AC power-flow optimization. MFH EV demand inputs are provided by ElectroTempo as a geolocated Level-2 demand layer; for Feeder 5 the screened demand is 384.0 kW for planning year 2025 at $P = 7.2$ kW per charger. The four feeders are Feeder 5 (*SAN LUIS OBISPO 1108*, PG&E ID 182631108), Feeder 124 (*BARTON 1108*, PG&E ID 253571108), Feeder 2996 (*CONTRA COSTA 2116*, PG&E ID 13652116), and Feeder 2707 (*BAY MEADOWS 2101*, PG&E ID 24012101). The implementation also performs a bounded β sweep; the effective upper bound is $\beta = 0.2$, explaining why served-demand percentages can reach 120% but not exceed it. k-means is run separately for each feeder and scenario. Principal parameters are summarized in Table I.

C. MFH spatial structure and k-means clustering

Let M denote the set of MFH demand points in the feeder service area. Each point $m \in M$ is associated with planar coordinates x_m and a demand proxy d_m. To summarize the coarse spatial structure of MFH demand, we apply k-means clustering to the MFH coordinates on a per-feeder basis and obtain the centroid set $\{c_1, \ldots, c_K\}$. In the reported experiments, K is either fixed by scenario or evaluated over a small sensitivity grid, and all reported values of K are stated explicitly in the results tables. The centroids capture coarse MFH demand hotspots, while the original MFH points retain street-level structure.

D. Node-level distance features

For each feeder node $i \in V$ with coordinates x_i, we compute two distance features:

$$d_i^{\text{cent}} = \min_{k=1,\ldots,K} \operatorname{dist}(x_i, c_k),$$
$$d_i^{\text{MFH}} = \min_{m \in M} \operatorname{dist}(x_i, x_m). \tag{1}$$

To make values comparable across feeders, each distance field is normalized to $[0, 1]$:

$$\hat{d}_i^{\text{cent}} = \frac{d_i^{\text{cent}}}{\max_{j \in V} d_j^{\text{cent}}}, \qquad \hat{d}_i^{\text{MFH}} = \frac{d_i^{\text{MFH}}}{\max_{j \in V} d_j^{\text{MFH}}}. \tag{2}$$

We then combine them into a single node-level distance score,

$$\text{score}_i = \eta \hat{d}_i^{\text{MFH}} + (1 - \eta)\hat{d}_i^{\text{cent}}, \qquad 0 \le \eta \le 1, \tag{3}$$

where η controls the relative importance of proximity to individual MFH buildings versus proximity to broader MFH demand hotspots. Nodes close to both MFH buildings and centroids receive scores near zero, while peripheral nodes far from MFH demand receive scores near one.

This score is purely geometric; it does not depend on line capacities or flows. It is the key ingredient for shaping the ℓ_1 penalty.

E. From distance scores to site penalties

The node-level score is converted into an activation penalty through

$$\lambda_i = \lambda_0 \left(1 + \alpha \, \text{score}_i\right), \tag{4}$$

where $\lambda_0 > 0$ is the base penalty and $\alpha \ge 0$ controls how strongly remote nodes are disfavored relative to MFH-adjacent nodes. Because the model uses a continuous activation variable $z_i \in [0, 1]$, this term acts as a convex sparsity-inducing surrogate rather than a strictly binary site-opening cost. Larger penalties suppress low-value remote nodes and concentrate service on nodes that are both electrically usable and spatially close to MFH demand.

F. Feeder graph and MFH demand attachment

We now briefly describe the feeder representation and how MFH demand is mapped onto nodes.

Each feeder is represented as a radial graph $G = (V, E)$ obtained from cleaned line geometry. A root node $r \in V$ is identified at the substation, and all edges are oriented away from the root to form a directed tree. A virtual source node src is added with one edge (src, r), producing the extended edge set E^{ext}. Each directed edge $e \in E^{\text{ext}}$ is assigned a capacity C_e, interpreted here as a thermal or load-hosting-capacity screening limit.

MFH demand is attached to the feeder graph in three stages. First, each MFH point is projected to its nearest feeder segment. Points farther than a cutoff radius $D_{\max}$ are excluded from the feeder-specific demand set. Second, a baseline capacity-only solve is used to identify electrically relevant flow corridors under unconstrained spatial bias. Third, projected MFH points are attached to segment endpoints associated with those corridors, producing node-level attached demand values d_i. This flow-informed attachment is intended to avoid assigning demand to electrically disconnected or weakly relevant feeder tails. The general framework supports this projection-based attachment, but the experiments reported in this paper disable that stage and instead use the MFH

locations themselves as the candidate pool for clustering and distance evaluation. Consequently, $D_{\max}$ is inactive in the reported tables, and the node-level demand quantities should be interpreted as screened siting demand rather than as the output of the projection routine.

G. Unified KCL + ℓ_1 formulation

With λ_i and d_i defined, we pose a single linear program that covers both the capacity-only and penalized cases.

a) Decision variables and parameters.: Decision variables are: $f_e \geq 0$ (real power flow on edge e, kW); $s_i \geq 0$ (EV demand served at node i, kW); $z_i \in [0,1]$ (relaxed site activation variable). Key parameters are: C_e (kW), branch capacity; d_i (kW), MFH-based EV demand at node i; $\beta \geq 0$, demand headroom factor for service up to $(1+\beta)d_i$; LHC_i (kW), node-level upstream hosting capacity; $M_i = \min\{(1+\beta)d_i, LHC_i\}$, upper bound for s_i; λ_i, distance-based site penalty; $\delta_i^{\text{in}}, \delta_i^{\text{out}}$, sets of flow edges entering and leaving node i.

b) Optimization problem.: We solve

$$\max_{f,s,z} \quad \sum_{i\in\mathcal{V}} s_i \; - \; \sum_{i\in\mathcal{V}} \lambda_i z_i \tag{5}$$

$$\text{s.t.} \quad 0 \leq f_e \leq C_e, \qquad \forall e \in \mathcal{E}^{\text{ext}}, \tag{6}$$

$$0 \leq s_i \leq (1+\beta)d_i, \qquad \forall i \in \mathcal{V}, \tag{7}$$

$$0 \leq s_i \leq LHC_i, \qquad \forall i \in \mathcal{V}, \tag{8}$$

$$s_i \leq M_i z_i, \qquad \forall i \in \mathcal{V}, \tag{9}$$

$$\sum_{e\in\delta_i^{\text{in}}} f_e - \sum_{e\in\delta_i^{\text{out}}} f_e = s_i, \qquad \forall i \in \mathcal{V}, \tag{10}$$

$$0 \leq z_i \leq 1, \qquad \forall i \in \mathcal{V}. \tag{11}$$

Modeling note on constraint (10). The KCL constraint sets net inflow equal to EV-served demand s_i, which implicitly assumes that no background (residential or commercial) load is present at feeder nodes beyond the EV component. This is a deliberate simplification for a screening model: we are sizing the available feeder headroom for new EV load, not simulating full dispatch. Planners should account for existing base load separately when interpreting the branch-utilization outputs.

In this model, f_e represents real-power flow on branch e, s_i represents EV charging demand served at node i, and z_i is a relaxed activation variable that links node service to a sparsity penalty. The formulation preserves linearity by allowing z_i to vary continuously on $[0,1]$. As a result, the model should be interpreted as producing sparse service patterns rather than strict binary facility-opening decisions. Reported active sites are therefore obtained using a small post-solve threshold on served demand or activation level. A binary $z_i \in \{0,1\}$ variant would convert the model to a MILP and was not part of the present multi-feeder sweep; preserving the LP form is the central modeling choice of this work.

c) Baseline KCL as a degenerate ℓ_1 case.: The baseline KCL model is recovered by setting $\lambda_i = 0$ for all nodes and fixing $z_i = 1$, or equivalently by removing the penalty and gating structure. In the reported experiments, this baseline is GIS-first: candidate EVSE locations are fixed in advance from spatial MFH information before the feeder-constrained solve. Location does not enter the objective, so the KCL solve serves only as an electrical screening step over that fixed candidate set. The penalized formulation preserves all feeder constraints and adds only MFH-aware distance penalties, so any change in the final service pattern reflects integrated spatial preference and sparsification, not a change in the underlying network model.

H. Solution procedure

For each feeder-scenario pair, we: (1) preprocess feeder geometry into a directed radial graph with branch capacities C_e; (2) compute MFH centroids, node-level distance features, and penalty values λ_i; (3) solve the baseline capacity-only model to identify initial feeder flow corridors; (4) attach MFH demand to feeder nodes via the segment-projection and flow-informed attachment rule; (5) solve the unified model in baseline and penalized modes over the specified parameter grid; and (6) record served demand, active-site counts, MFH proximity metrics, branch utilization, bottlenecks, and runtime.

All optimization runs are carried out in JuMP with the HiGHS solver. Solver version, stopping criteria, hardware, and runtime settings are reported so that results can be replicated.

I. Experimental configuration

The study uses a four-feeder comparison together with a detailed parameter sweep on Feeder 5. Across all feeders, the EV charger size was fixed at $P = 7.2$ kW (Level 2), the base penalty weight was set to $\lambda_0 = 7.2$, and the MFH/centroid distance blend was fixed at $\eta = 0.6$, corresponding to a 60% MFH-distance contribution and a 40% centroid-distance contribution. The implementation also performs an internal β sweep over relaxed demand levels; the effective upper bound observed in the reported results is $\beta = 0.2$. For the L1 runs, results were averaged over four random seeds: 42346, 42347, 42350, and 42352. MFH attachment was disabled in this experiment. Timing was recorded using the `@elapsed` macro after a JIT warmup pre-run.

The feeder set consists of feeders 5, 124, 2996, and 2707, with screened EV demands of 384.0, 710.4, 1612.8, and 3244.8 kW, respectively. The detailed map-based illustrations in Figs. 1 and 2 use Feeder 5. For Feeder 5 and Feeder 2707, the centroid counts tested were $K \in \{1,3,5,7\}$; for Feeders 124 and 2996, the centroid counts were expanded to $K \in \{1,3,5,7,10\}$ because those runs were executed with a larger centroid grid. Across all cases, the penalty multipliers were swept over $\alpha \in \{2,5,10,18,25\}$. In the multi-feeder summary, the reported "best" L1 configuration for each feeder is the one with the highest EV capacity served, with ties broken by the lowest math-core runtime.

III. Results

This section evaluates how MFH-biased distance penalties affect feeder-level EV demand screening, spatial siting patterns, runtime, and branch-utilization behavior in the unified

TABLE II
FEEDER 5 SCREENING RESULTS. ALL L1 CONFIGURATIONS REDUCE MEAN MFH DISTANCE FROM 411.6 M (KCL BASELINE) TO 271.0 M, DEMONSTRATING THAT THE ℓ_1 PENALTY SHIFTS SITES TOWARD MFH DEMAND WITHOUT REDUCING FEEDER SERVICEABILITY.

Method	K	α	Served [kW]	**Served [%]**	Mean MFH [m]	Time [s]
KCL	–	–	384.0	**100.0**	411.6	0.0315
L1	1	18	422.4	**110.0**	271.0	0.2089
L1	3	25	403.2	**105.0**	271.0	0.1982
L1	5	25	403.2	**105.0**	271.0	0.1969
L1	7	25	441.6	**115.0**	271.0	0.1632

KCL+ℓ_1 model of (5)–(11). The baseline KCL configuration is treated as the no-penalty reference case, obtained by setting $\lambda_i = 0$ and fixing $z_i \equiv 1$. In the reported experiments, that baseline functions as a GIS-first electrical-screening case: MFH-informed candidate locations are fixed in advance from spatial data, and the feeder-constrained solve determines only how much of the prescribed demand can be served over that fixed candidate set. Penalized cases preserve all feeder constraints and differ only through the MFH-aware activation penalties.

A. Screening objective and evaluation metrics

The primary metric is total served load $S_{\text{tot}} = \sum_{i \in V} s_i$, interpreted as a binary screening quantity (can the feeder serve the target demand?). Served-demand percentage $= 100 \times \sum_i s_i / \sum_i d_i$ is reported as a secondary metric; values above 100% reflect the optimizer using β-relaxed headroom, not a constraint violation. Spatial quality is assessed via mean MFH distance

$$\bar{d}_{\text{MFH}} = \frac{1}{|A|} \sum_{i \in A} \min_{m \in M} \text{dist}(x_i, x_m), \tag{12}$$

where $A = \{i \in V : z_i > 0.001\}$ is the active-site set. Branch utilization $u_e = f_e / C_e$; a branch is flagged as a bottleneck when $u_e \geq u^{\text{th}}$ (e.g., 0.8 or 0.9). Bottleneck frequency is the fraction of evaluated configurations in which a branch is flagged. Timing uses math-core-only time after JIT warmup, excluding graph I/O and plotting.

B. Feeder 5 screening results

Table II summarizes representative Feeder 5 results. The question is binary: can the screened EV demand of 384.0 kW be served under the feeder-aware LP constraints? The KCL baseline answers it for a fixed GIS-derived candidate set; the L1 runs keep that screening role but reshape *where* service lands through MFH-biased penalties.

The baseline KCL model serves the full 384.0 kW target (100.0%), establishing that the feeder can accommodate the screened demand for the fixed candidate set. The L1 runs are therefore spatial refinements: their main value is reducing mean MFH distance from 411.6 m to 271.0 m while preserving feeder serviceability. The $K = 7$, $\alpha = 25$ configuration achieves the highest served capacity at 441.6 kW (115%), as finer centroid resolution more precisely distinguishes MFH-adjacent nodes from peripheral ones, allowing the optimizer to concentrate flow on compact, electrically favorable, MFH-proximate corridors.

TABLE III
FOUR-FEEDER SUMMARY: SERVED DEMAND UNDER BASELINE KCL AND BEST L1.

Feeder	Nodes	n_{MFH}	Demand [kW]	KCL [%]	L1 [%]
5	350	16	384.0	100.0	120.0
124	329	32	710.4	52.1	52.1
2996	269	35	1612.8	100.6	104.7
2707	372	82	3244.8	104.8	110.0

TABLE IV
FOUR-FEEDER SUMMARY: SPATIAL METRICS AND BOTTLENECK BEHAVIOR.

Feeder	Mean MFH [m]	Median [m]	Min–Max [m]	Bneck. [%]
5	365.3	424.2	64.0–712.7	0
124	99.2	76.3	43.1–373.3	100
2996	312.6	199.9	28.0–1033.6	0
2707	81.6	55.1	13.1–255.3	0

C. Multi-feeder screening comparison

Tables III and IV extend the analysis to all four feeders. "Best" L1 means highest EV capacity served with ties broken by lowest math-core runtime; bottleneck frequency is the fraction of evaluated $(K \times \alpha \times \text{seed})$ configurations for which the feeder fails to serve 100% of screened MFH demand.

Three capacity regimes emerge. Feeder 5 is *headroom-rich*: the baseline serves the full target and L1 improves MFH proximity on top. Feeder 124 is *hard-constrained*: both KCL and the best L1 plateau at **52.1%**; binding branch limits cap performance regardless of penalty tuning, and the bottleneck frequency is 100%—no amount of tuning recovers capacity the topology lacks. Feeders 2996 and 2707 are *partially-constrained*: both exceed their screened demand thresholds, but Feeder 2707 achieves far stronger MFH proximity (mean 81.6 m vs. 312.6 m), showing that feeder topology and MFH geography jointly shape the spatial payoff of the penalized formulation.

D. Spatial effect of MFH-biased penalties

Figures 1 and 2 show the spatial effect of adding MFH-biased distance penalties to the objective. In the baseline GIS-first KCL case, the candidate set is fixed from spatial inputs and the model allocates feasible EV service wherever branch headroom is available. That produces a diffuse pattern with many lightly used sites, including nodes that are electrically serviceable but weakly aligned with MFH demand. Once MFH-aware penalties are activated, the solution suppresses remote, low-value sites and concentrates service on nodes that are both electrically usable and geographically closer to MFH concentrations.

The proximity metrics tell the same story. The baseline KCL run yields a mean MFH distance of 411.6 m. Several L1 settings reduce that to 271.0 m while still serving at least the screened demand threshold. The ℓ_1 term leaves the feeder constraints intact — what shifts is which feasible nodes the optimizer favors.

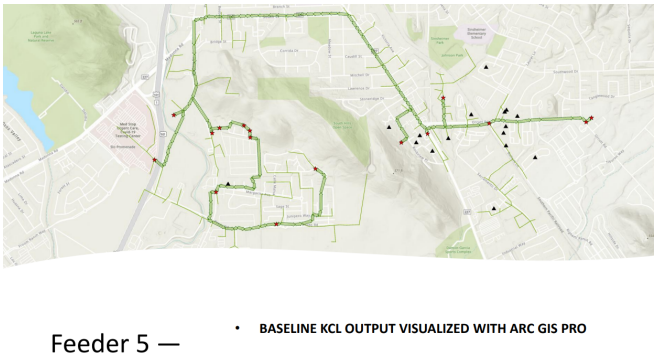

Fig. 1. Baseline GIS-first KCL screening results on Feeder 5 (16 MFH sites, 16 active sites). Sites are spatially diffuse because no MFH-proximity penalty is applied; the optimizer places service wherever branch headroom is available, including electrically accessible but MFH-distant nodes. Mean MFH distance: 411.6 m. Map prepared by the Texas A&M group from PG&E Distribution Capacity / GRIP GIS source layers [19].

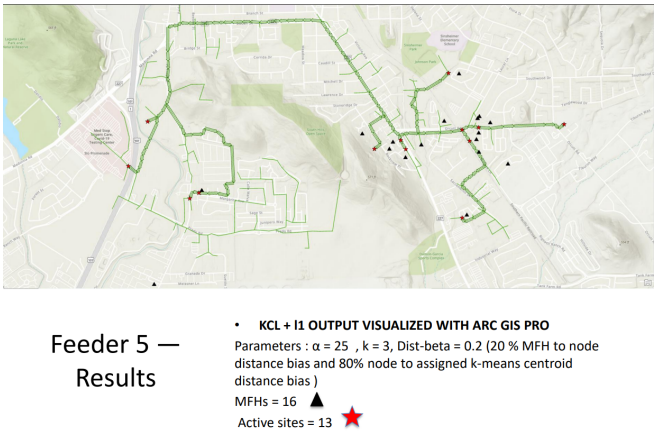

Fig. 2. KCL+ℓ_1 siting results on Feeder 5 ($K = 3$, $\alpha = 25$; 16 MFH sites, 13 active sites). Compared with Fig. 1, the penalty suppresses peripheral sites and concentrates service on MFH-adjacent nodes, reducing mean MFH distance from 411.6 m to 271.0 m while preserving at least the screened demand level. Map prepared by the Texas A&M group from PG&E Distribution Capacity / GRIP GIS source layers [19].

E. Runtime and parameter-sweep behavior

The screening model is inexpensive to solve. Both KCL and L1 solve times remain in the sub-second range after JIT warmup (Table V): KCL math-core runtime spans 0.0068–0.0932 s; best-L1 runtime spans 0.0109–0.1505 s. Runtime is not determined by node count alone—Feeder 2707 (372 nodes) runs faster than Feeder 2996 (269 nodes) at the best-L1 configuration—reflecting dependence on MFH candidate count and the tested K set as well as feeder size.

The sweep also reveals a practical tradeoff: lower-to-moderate penalty settings ($\alpha = 18$, $K = 1$) achieve 110% served capacity on Feeder 5, while higher settings ($\alpha = 25$, $K \geq 3$) achieve the same MFH proximity of 271.0 m but differ in served capacity by centroid resolution. Planners who prioritize MFH proximity should select higher α values; those who prioritize maximizing served capacity may prefer moderate α settings.

TABLE V
FOUR-FEEDER SUMMARY: BEST L1 CONFIGURATION AND RUNTIME.

Feeder	Best K	Best α	Sites placed	L1 math [s]	Avg. L1 math [s]
5	7	10	12	0.0372	0.0519
124	1	10	15	0.0109	0.0220
2996	7	25	9	0.1505	0.1478
2707	7	25	61	0.0263	0.0212

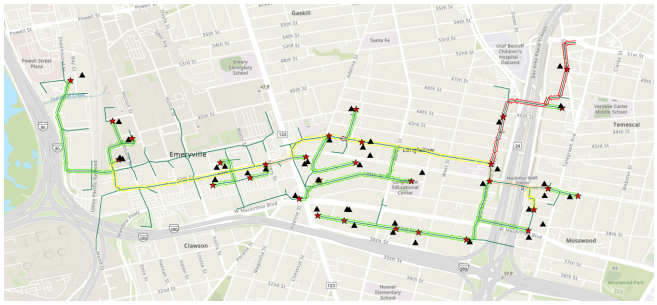

Fig. 3. Branch-utilization map for an Oakland/Emeryville feeder. Dark red segments are near capacity limits; green segments have low utilization. Map from PG&E Distribution Capacity / GRIP GIS source layers [19].

F. Binary-z versus continuous-z formulation

We examined whether replacing $z_i \in [0, 1]$ with a binary $z_i \in \{0, 1\}$ would materially change results. Across 420 evaluated feeder-configuration instances, Feeders 5 and 2707 produced fully integral z values in every configuration; Feeders 124 and 2996 showed only limited fractional activity at a small number of nodes with negligible impact on delivered power. The LP relaxation therefore has a negligible practical integrality gap, capturing the relevant siting behavior at the speed of a linear program—no integer branching required. A broader mixed-integer comparison is left for future work.

G. Feeder loading and bottleneck interpretation

Every configuration enforces $0 \leq f_e \leq C_e$, so no solution violates the prescribed feeder screening limits. In baseline mode, flow spreads across many branches including feeder tails. Under MFH-biased penalties, flow concentrates on corridors connecting MFH-adjacent sites back toward the source; the utilization distribution grows more skewed and a small set of branches recurs repeatedly as heavily used EV-serving corridors.

Figure 3 illustrates this: branches interpreted as candidate bottlenecks repeatedly appear near their capacity limits across MFH-biased siting solutions and surface naturally from the LP—not from ad hoc heuristics. Heavily loaded segments in Feeder 2707 show utilization values of 50–75%, while lightly loaded tails run below 50%.

Because these bottleneck indicators emerge from the same LP that produces site recommendations, they are directly interpretable. Segments are not flagged by ad hoc heuristics. They surface because they repeatedly appear as highly utilized links across MFH-aligned screening solutions. Dark red segments in the utilization visualization carry the highest loading; green segments indicate spare capacity.

H. Planning implications and limitations

The feeder-aware LP answers two distinct questions for a utility planner: first, can this feeder serve the target EV load under thermal and KCL constraints?; and second, which candidate sites align with concentrated MFH demand? Feeder 124 makes the first question binding—no penalty tuning recovers capacity the topology lacks. Feeders 5, 2996, and 2707 move on to the second question, with Feeder 2707 showing the strongest MFH pull when headroom is available. Candidate sites close to MFH demand while drawing on lightly loaded upstream corridors are strong no-regrets options; repeatedly stressed branches pinpoint where conductor upgrades or reconductoring would unlock additional EV access.

Several extensions remain outside the current scope. Binary-z_i sensitivity testing and downstream AC power-flow validation are the most immediate next steps. Neither changes the main finding—the LP reliably screens feeder-level demand serviceability and prioritizes EV siting nodes across multiple feeders—but both define where the current contribution ends and future work begins.

IV. Conclusion

Feeder-aware linear programming can serve two distinct planning functions at once: screening whether a distribution feeder can support a prescribed EV demand level under thermal and KCL constraints, and spatially prioritizing candidate sites toward concentrated MFH demand. Neither function requires integer variables or nonconvex constraints. The LP stays light enough to rerun across many feeders and parameter combinations without solver tuning.

The most practically significant result is the behavior of the four-feeder comparison. Feeder 124 demonstrates that capacity limits can be binding enough to cap served demand regardless of how the penalty is tuned—a finding a planner cannot easily recover from spatial heuristics alone. By contrast, Feeders 2707 and 2996 show that where headroom exists, the ℓ_1 penalty meaningfully shifts siting outcomes: mean MFH distances drop from 312–366 m to as low as 81 m without sacrificing serviceability. The binary-versus-continuous sensitivity analysis further shows that the LP relaxation captures binary siting behavior on these feeders with negligible integrality gap, supporting the use of the faster continuous formulation in production screening.

Beyond siting, the branch-utilization outputs provide bottleneck indicators directly useful for utility planning. Recurrently stressed segments emerge naturally from the LP as candidate upgrade locations, linking demand screening, spatial prioritization, and early-stage reinforcement guidance within a single optimization layer.

These results carry a clear scope limitation. The model approximates feeder thermal capacity and omits full AC power-flow physics. Future work should target time-varying charging demand, larger cross-feeder validation, binary-versus-relaxed activation sensitivity, and downstream AC validation on higher-fidelity feeder models. Detailed electrical-parameter optimization and upgrade modeling can then be applied selectively where the screening layer identifies the highest planning value.

References

[1] T. Unterluggauer, J. Rich, P. B. Andersen, and S. Hashemi, "Electric vehicle charging infrastructure planning for integrated transportation and power distribution networks: A review," *eTransportation*, vol. 12, p. 100163, 2022.

[2] R. Pagany, L. Ramirez Camargo, and W. Dorner, "A review of spatial localization methodologies for the electric vehicle charging infrastructure," *International Journal of Sustainable Transportation*, vol. 13, no. 6, pp. 433–449, 2019.

[3] L. Hanig, C. Ledna, D. Nock, C. D. Harper, A. Yip, E. Wood, and C. A. Spurlock, "Finding gaps in the national electric vehicle charging station coverage of the United States," *Nature Communications*, vol. 15, no. 1, p. 4871, 2024.

[4] Q. Ren and M. Sun, "Predicting the spatial demand for public charging stations for EVs using multi-source big data: An example from Jinan City, China," *Scientific Reports*, vol. 15, no. 1, p. 6991, 2025.

[5] G. Raman, G. Raman, and J. C.-H. Peng, "Resilience of urban public electric vehicle charging infrastructure to flooding," *Nature Communications*, vol. 13, no. 1, p. 3213, 2022.

[6] M. Kuby and C. Upchurch, "The flow-refueling location problem for alternative-fuel vehicles," *Transportation Research Part B: Methodological*, vol. 38, no. 7, pp. 593–620, 2004.

[7] I. Capar and M. Kuby, "An efficient formulation of the flow refueling location model for alternative-fuel stations," *IIE Transactions*, vol. 44, no. 8, pp. 622–636, 2012.

[8] A. Y. S. Lam, Y.-W. Leung, and X. Chu, "Electric vehicle charging station placement: Formulation, complexity, and solutions," *IEEE Transactions on Smart Grid*, vol. 5, no. 6, pp. 2846–2856, 2014.

[9] Y. Zhang, Q. Zhang, A. Farnoosh, S. Chen, and Y. Li, "GIS-based multi-objective particle swarm optimization of charging stations for electric vehicles," *Energy*, vol. 169, pp. 844–853, 2019.

[10] H. Zhang, S. J. Moura, Z. Hu, W. Qi, and Y. Song, "A second-order cone programming model for planning PEV fast-charging stations," *IEEE Transactions on Power Systems*, vol. 33, no. 3, pp. 2763–2777, 2018.

[11] P. Sadeghi-Barzani, A. Rajabi-Ghahnavieh, and H. Kazemi-Karegar, "Optimal fast charging station placing and sizing," *Applied Energy*, vol. 125, pp. 289–299, 2014.

[12] S. Deb, K. Tammi, X.-Z. Gao, K. Kalita, P. Mahanta, and S. Cross, "A robust two-stage planning model for the charging station placement problem considering road traffic uncertainty," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 6571–6585, 2022.

[13] M. R. Abdullahi, Q.-C. Lu, A. Hussain, S. Tripura, P.-C. Xu, and S. Wang, "Location optimization of EV charging stations: A custom k-means cluster algorithm approach," *Energy Reports*, vol. 12, pp. 5367–5382, 2024.

[14] E. Magsino, F. M. M. Espiritu, and K. D. Go, "Discovering electric vehicle charging locations based on clustering techniques applied to vehicular mobility datasets," *ISPRS International Journal of Geo-Information*, vol. 13, no. 10, p. 368, 2024.

[15] P. Geng *et al.*, "Intelligent layout and optimization of EV charging stations: Initial configuration via enhanced k-means and subsequent refinement through integrated GCN," *Radioengineering*, vol. 34, no. 1, pp. 79–91, 2025.

[16] J. Z. Khameneh and E. Fagbenle, "Optimal placement of electric vehicle chargers: A mixed-integer linear programming model," *Clean Technologies and Recycling*, vol. 6, no. 1, pp. 1–32, 2026, doi: 10.3934/ctr.2026001.

[17] B. Palmintier, T. Elgindy, C. Mateo, F. Postigo, T. Gomez, F. de Cuadra, and P. Martinez, "Experiences developing large-scale synthetic U.S.-style distribution test systems," in *Proc. 21st Power Systems Computation Conf. (PSCC)*, 2020, doi: 10.1016/j.epsr.2020.106665.

[18] P. Paudyal, S. Ghosh, S. Veda, D. Tiwari, and J. Desai, "EV hosting capacity analysis on distribution grids," in *Proc. IEEE Power and Energy Society General Meeting (PESGM)*, 2021, doi: 10.1109/PESGM46819.2021.9638170.

[19] PG&E, "Distributed Resource Planning / Grid Resource Integration Portal," [Online]. Available: https://www.pge.com/en/business/environment/drp.html (accessed Apr. 2026).

[20] PG&E, “Grid Resource Integration Portal quick start guide,” [Online]. Available: https://www.pge.com/en/business/environment/drp/grip-guide.html (accessed Apr. 2026).