V2I-Coordinated Fair Power Allocation for Electric Vehicles on Dynamic Charging Lanes

Gauthier Le Tat
LIGM CNRS-UMR 8049
University Gustave Eiffel
Champs-sur-Marne, France
gauthier.le-tat@univ-eiffel.fr

Ilyes Ahmim
LIGM CNRS-UMR 8049
University Gustave Eiffel
Champs-sur-Marne, France
ilyes.ahmim2@univ-eiffel.fr

Abderrezak Rachedi
LIGM CNRS-UMR 8049
University Gustave Eiffel
Champs-sur-Marne, France
abderrezak.rachedi@univ-eiffel.fr

Mohamed Chachoua
LASTIG
EIVP, University Gustave Eiffel
Paris, France
chachoua@eivp-paris.fr

Ahmed Mebarki
MSME UMR 8208
University Gustave Eiffel
Champs-sur-Marne, France
ahmed.mebarki@univ-eiffel.fr

Fabrice Vienne
COSYS/PICS-L
University Gustave Eiffel
Champs-sur-Marne, France
fabrice.vienne@univ-eiffel.fr

***Abstract*—Dynamic Charging Lanes (DCLs) recharge electric vehicles in motion, coordinated by a real-time control loop linking vehicles to a roadside scheduler through vehicle-to-infrastructure (V2I) communications. At high traffic densities, the lane may operate in permanent saturation: the aggregate charging demand would exceed the grid power budget. However, the only existing real-time formulation of this problem assumes that every vehicle can be served, thus breaking down precisely in this regime. This paper models the DCL as a networked control system whose event-driven loop composed of IEEE 802.11bd V2I and wired links, and reformulates power allocation as a bi-objective problem, combining a weighted global utility with a weighted Jain's fairness index and tackling the saturated regime issue. Two scheduling heuristics are proposed, a weighted max-min policy, and a utility/fairness mixed family, both evaluated against two single-objective baselines. Event-based simulations, with realistic parameters from European vehicles' data, show that the mixed family traces a tunable Pareto frontier between fairness and utility and dominates the weighted max-min algorithm, while computing each allocation in about one millisecond. Its O(nlog n) complexity ensures both scalability and real-time viability, while a Jain's fairness index close to 0.95 in realistic scenarios confirms the equity of the resulting allocation, demonstrating the practicality of fair DCL scheduling in the saturated regime.**

***Index Terms*—V2I communications, power allocation, fairness, dynamic charging lane, IEEE 802.11bd, real-time control**

I. Introduction

The rapid evolution of intelligent transportation systems is reshaping mobility by integrating sensing, communication, and control technologies to improve traffic efficiency, reliability, and safety. Among these technologies, electric vehicles (EVs) help cut greenhouse-gas emissions and fossil-fuel dependency. However, they remain held back by basic battery limitations (limited range, charging time, battery weight, infrastructure availability), particularly in heavy-duty transport such as freight trucks and long-haul logistics, where high vehicle mass drives much higher energy consumption.

Electric Road Systems (ERS) have been proposed as effective solutions to these issues, and Connolly [1] estimates them cheaper than heavier batteries or oil dependence. Various dynamic charging technologies (inductive wireless power transfer, conductive rail, or overhead catenary) all aim to enable recharging in motion, reducing reliance on stationary infrastructure and range anxiety [2]. In Dynamic Charging Lane (DCL) systems, a centralized controller coordinates power allocation among vehicles over vehicle-to-infrastructure (V2I) communication, scheduling charging in real time as vehicles cross the lane.

Despite their potential, DCLs introduce significant power-distribution challenges. Vehicles remain in the lane only briefly, requiring high transfer rates, yet the available power is inherently limited by grid supply and infrastructure capacity. When traffic density is high, the aggregated demand exceeds what can be supplied simultaneously, so vehicles compete for a constrained power budget. Furthermore, this demand is highly heterogeneous (depending on vehicle type, battery capacity, and load, with heavy-duty vehicles using far more power than cars), so allocation strategies must fairly balance diverse requirements while keeping infrastructure utilization efficient.

This paper addresses power distribution scheduling in DCLs under permanent saturation, where the infrastructure delivers its maximum power yet cannot satisfy all vehicles at once. While power allocation in static charging has been studied [3], [4], real-time scheduling within a single DCL has received little attention so far. Qiu et al. [5] addressed it under a feasibility assumption (every vehicle leaves with enough energy to reach its destination) that no longer holds as traffic density grows. This paper therefore makes three contributions. First, we reformulate the allocation problem under persistent saturation as a well-posed bi-objective problem trading global utility (the weighted number of vehicles served) against fairness: prioritizing urgent vehicles lowers fairness, whereas strictly fair sharing leaves more unsatisfied demands. Second, we propose a tunable utility/fairness mixed family

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

of computationally efficient $O(n \log n)$ heuristics suitable for real-time deployment, alongside a weighted max-min policy as a natural fairness-oriented candidate. Third, modeling the DCL as a networked control system whose event-driven loop runs over IEEE 802.11bd V2I links, we characterize its end-to-end latency and show that the loop closes within about 11 ms (under 16 cm of vehicle travel), confirming the real-time viability of fair DCL scheduling under saturation. This regime, largely unexplored to date, could in the longer term become a practical concern as electric roads are deployed more widely.

The remainder of this paper is organized as follows. Section II reviews related work; Section III presents the system model and problem formulation; Section IV details the scheduling algorithms (two baselines and two proposed heuristics); Section V reports the simulation evaluation; and Section VI concludes.

II. Related Work

A. Charging Demand Management at the Network Level

Recent research on EV charging coordination has primarily focused on large-scale interactions between transportation networks and power distribution systems, jointly optimizing vehicle routing, charging-infrastructure placement, and electricity pricing through equilibrium-based models of coupled transportation and power dynamics. Li et al., for instance, investigate co-optimized operation under mixed user equilibrium, where EV drivers choose routes and charging strategies while accounting for both travel cost and electricity pricing [6].

B. Dynamic Wireless Charging Infrastructure

A second line of work studies the planning and deployment of dynamic charging infrastructure. Liu et al. jointly determine the locations and electricity prices of dynamic charging links through a bi-level model in which drivers adapt their routes and recharging plans, so as to minimize the total social cost under a given budget [7]. Zhang and Yu couple infrastructure planning with grid operation in a two-stage framework that places the charging tracks and then coordinates the charging and discharging of the EVs over them for vehicle-to-grid services [8]. More recently, Wang et al. propose a multi-stage distributionally robust framework that jointly optimizes lane deployment and energy allocation under traffic and renewable uncertainties, balancing cost, emissions, and grid resilience [9]. These frameworks, however, decide at the city scale and address infrastructure planning, not the real-time, instantaneous distribution of power among vehicles sharing one lane.

C. Allocation under Limited Charging Facilities

A complementary line of research addresses charging congestion from *limited charging facilities*, closely related to the saturated regime studied here. Bai et al. propose a distributed coordination framework for electric trucks competing for a limited number of charging ports at stationary stations, each truck planning from waiting-time estimates under travel and energy-consumption uncertainties [10]. This shares our central premise that charging resources cannot satisfy all demands at once, but tackles it through decentralized, first-come-first-served port assignment over discrete charging stops and routing horizons, rather than the continuous, event-driven sharing of a power budget among vehicles co-located on one lane.

In a second stream, the limited resource is a continuous *power* budget shared at a capacity-constrained coupling point. Fridgen et al. study residential charging under grid bottlenecks and allocate power by each driver's *travel time*, defined as the driving time plus the extra time of an external charging stop (detour plus recharging) whenever the departure state of charge (SOC) is too low to reach the destination [4]. The key feature, which inspired our notion of service, is that this cost is fundamentally *non-linear*: negligible while the stored energy suffices, it jumps once an extra stop becomes necessary, motivating a *thresholded* measure of service rather than one proportional to delivered energy. Their setting is nonetheless static residential charging, where charging can be time-shifted; we transpose this travel-time-based service to the dynamic lane, where vehicles share the infrastructure only briefly and power must be reallocated continuously.

D. Real-Time Power Allocation in a Single Charging Lane

To the best of our knowledge, Qiu et al. [5] is the only study that explicitly addresses real-time power allocation among multiple EVs within a single dynamic wireless charging lane: each entering vehicle reports its SOC, position, and velocity to a central controller that schedules the power delivered by each section it crosses, within the lane's capacity. Despite its relevance, the formulation still enforces per-vehicle feasibility through hard constraints (in particular, that every vehicle leaves with enough energy to reach its destination), which becomes unrealistic under limited supply or high traffic density, where trade-offs are unavoidable. In contrast, our formulation removes this strong constraint and studies scheduling under persistent power shortage while relying on a similar system model, combining the objectives into a unified bi-objective formulation balancing fairness and global utility, a more realistic model of heterogeneous demands competing for limited resources.

E. Real-Time V2I Communication for Closed-Loop Control

Closing the allocation loop in real time relies on low-latency V2I exchanges, whose performance has been characterized independently of charging. Charpentier et al. report field latencies between the on-board unit (OBU) and road-side unit (RSU) over ITS-G5 [11], Coll-Perales et al. decompose end-to-end vehicle-to-everything (V2X) latency in 5G networks [12], and Xue et al. show that IEEE 802.11bd roughly halves the latency of 802.11p while improving reliability [13]. We build on these results to budget our control loop (Section V-D).

III. System Model and Problem Formulation

A. System Architecture

The system (Fig. 1) is a single DCL providing continuous wireless or conductive power transfer to EVs in motion. The

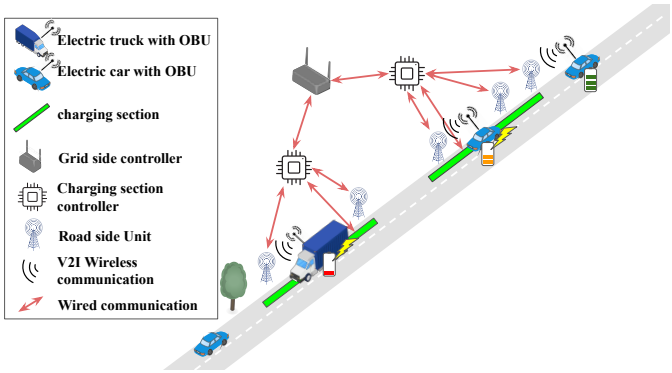


Fig. 1: System model of the dynamic charging lane.

lane carries N_s charging sections $[s_j^{\text{start}}, s_j^{\text{end}}]$, each one being an independent power transfer unit energizing the vehicles above it; they need be neither contiguous nor equally spaced. Beyond this physical layer, the system is a real-time networked control system with two interacting layers.

Power-delivery layer. Each section is driven by a Charging Section Controller (CSC), the power electronics converting grid energy into the field delivered to the vehicles. The CSCs take their power from a common supply whose aggregate budget A is set by the grid and managed by the Grid-Side Controller (GSC). On the vehicle side, an onboard receiver feeds the battery via the Battery Management System.

Communication and control layer. The GSC executes the scheduling algorithm on a roadside edge server, which in a wider deployment may host it alongside other electric-road services [14]; here we consider the scheduler in isolation. Each vehicle carries an OBU for V2I connectivity, while RSUs, spaced every few hundred meters for continuous short-range coverage, link the OBUs to the GSC over a wired backhaul; a wired control link connects the GSC to the CSCs.

The GSC thus closes a real-time control loop: it senses the vehicles, computes an allocation, and actuates the CSCs. The allocation is event-driven, recomputed whenever the state changes (a vehicle enters or exits a section, reaches full charge, or has its demand fulfilled), so the loop continuously adapts; its latency is analyzed in Section V-D.

B. Communication Model

Operating the DCL requires V2I exchanges (between vehicles and RSUs) and intra-infrastructure exchanges (among fixed entities), summarized in Table I.

1) V2I Communications: In the uplink, an entering vehicle issues a *registration* declaring its identifier and static profile (Table I); the operator maps this identifier to a vehicle class, from which the priority weight w_i follows (Section V-A); the vehicle does not choose it. While on the lane it periodically sends a *state report* of its position and SOC, so that the GSC tracks its presence and detects reallocation events. In the downlink, the GSC returns an *acknowledgment* and, after each reallocation, a *power grant* carrying the new authorized power $P_i(t)$. A secure authentication step is required but out of scope; lightweight protocols for similar resource-constrained settings confirm its feasibility [15], [16].

2) Intra-Infrastructure Communications: The RSUs relay uplink reports and downlink allocations over the wired backhaul, while the GSC drives the CSCs over a real-time control link (Table I). How a CSC physically directs power to each vehicle (energizing the right coil or rail segment) is technology-dependent and out of scope; we assume it delivers the assigned power, combining the GSC allocations with its own sensing, with a vehicle's session handed over between CSCs at section boundaries.

3) Technology and Deployment Assumptions: The model does not depend on the radio technology: any link delivering these messages within a few milliseconds (including 5G NR V2X) would suit the loop. We assume an infrastructure-centric deployment over IEEE 802.11bd, the recent V2X evolution of 802.11p [17]: operating in the 5.9 GHz ITS band in ad-hoc (OCB) mode, it lets the operator own the radio access while staying backward-compatible, roughly halving latency and improving reliability by about 20% [13]. As the vehicle speed v is constant and the SOC evolves deterministically from the granted power, the GSC predicts positions and energies between reports; the regular reports (every few seconds) thus mainly resynchronize the model and keep the channel lightly loaded (short range limits contention to only a few tens of vehicles per RSU).

TABLE I: Main communication messages and the model variables they carry.

Message (direction)	**Content**	**Trigger**
Registration (V $\rightarrow$ I)	$\text{ID}_i, B_i, c_i, P_i^{\max}, d_i^{\text{post}}$	entry
State report (V $\rightarrow$ I)	$p_i(t)$, $y_i(t)$	periodic
Section status (CSC $\rightarrow$ GSC)	delivered power	periodic
Power grant (I $\rightarrow$ V)	$P_i(t)$	per event
Section cmd. (GSC $\rightarrow$ CSC)	$P_i(t)$, $x_j(t)$	per event

C. Charging and Mobility Model

EVs continuously enter and exit the infrastructure; arrivals follow a Poisson process to capture traffic variability. On entry, each vehicle samples parameters from predefined distributions representing heterogeneity and crosses the lane at constant speed v. The GSC distributes the available power among all vehicles on the lane, recomputing the event-driven allocation at the state changes listed above, which keeps utilization efficient under changing traffic.

The main variables / parameters are summarized in Table II.

D. System Constraints

Based on the defined parameters, the operation of the dynamic charging lane is governed by the following system constraints.

TABLE II: Model Parameters and Variables

Symbol	Unit	Definition
Global		
λ	veh/s	Average vehicle arrival rate
A	kW	Total power of the charging infrastructure
L	km	Total length of the charging lane
Charging sections		
N_s	—	Number of charging sections
s_j^{start}, s_j^{end}	km	Start / end positions of charging section j
a_j	kW	Maximum power that can be delivered by charging section j
$x_j(t)$	kW	Power supplied by charging section j at time t
Vehicle-specific		
$\mathcal{V}$	—	Set of all vehicles (the exit population)
$\mathcal{V}(t)$	—	Vehicles on the lane at time t, $\mathcal{V}(t) \subseteq \mathcal{V}$; indexed symbols (e.g., $P_i(t)$) refer to the i-th vehicle
v	km/h	Vehicle speed on the charging lane
$y_i(t)$	—	SOC of vehicle i at time t
$p_i(t)$	km	Position of vehicle i at time t
$t_{i,e}$	s	Time at which vehicle i exits the lane
$y_{th,i}(t)$	—	Minimum SOC required for vehicle i to safely reach the next charging section
B_i	kWh	Battery capacity of vehicle i
c_i	kW	Traction power consumption of vehicle i
w_i	—	Priority weight associated with vehicle i
$P_i(t)$	kW	Current power allocated to vehicle i
$P_i^{\max}$	kW	Maximum charging power of vehicle i
d_i^{post}	km	Distance that vehicle i must travel after leaving the charging infrastructure

$$\sum_{j=1}^{N_s} x_j(t) = A, \qquad \forall t \quad (1)$$

$$x_j(t) \le a_j, \qquad \forall j,\ \forall t \quad (2)$$

$$y_i(t) \ge y_{th,i}(t), \qquad \forall i,\ \forall t \quad (3)$$

$$y_i(t) \le 1, \qquad \forall i,\ \forall t \quad (4)$$

$$P_i(t) \le P_i^{\max}, \qquad \forall i,\ \forall t \quad (5)$$

$$y_i(t+\Delta t) = y_i(t) + \frac{P_i(t) - c_i}{B_i}\Delta t, \qquad \forall i,\ \forall t \quad (6)$$

E. Objective Functions

The allocation problem optimizes two complementary objectives: a global utility formalizing the completion of individual energy demands, and the equity of power distribution among vehicles present on the lane.

1) Global Utility Objective: The first objective evaluates whether vehicles exit with enough energy to reach their destination, the next point where they can recharge (final endpoint or intermediate facility), so d_i^{post} is vehicle i's own distance from the lane exit to that target. As discussed in Section II-C, a vehicle's effective service is non-linear in its final energy [4]: unaffected while the stored energy suffices, it jumps once an extra charging stop becomes necessary.

This observation motivates a thresholded, non-linear formulation of the global utility. Rather than rewarding incremental energy gains, we define a binary utility function that returns 1 if the vehicle's stored energy at the exit of the lane meets the destination requirement and 0 otherwise:

$$U_i = \mathbf{1}\left(y_i(t_{i,e})B_i \ge \frac{d_i^{\text{post}} c_i}{v}\right) \quad (7)$$

where $\mathbf{1}(\cdot)$ is the indicator function and $\mathcal{V}$ the set of all vehicles. The Global Utility (GU) Objective is their weight-normalized sum, accounting for priority:

$$\max \frac{\sum_{i\in\mathcal{V}} U_i w_i}{\sum_{i\in\mathcal{V}} w_i}. \quad (8)$$

This yields a value in $[0, 1]$: the fraction of total vehicle-weight satisfied at the lane exit.

2) Fairness Objective: The global utility does not constrain how power is distributed at a given instant. The second objective therefore promotes fairness via Jain's fairness index [18], a standard equity measure recently adopted for real-time EV allocation under limited capacity, including with user prioritization [19]. Adapted to the priority weights, with $\mathcal{V}(t)$ the set of vehicles on the lane at time t (Table II), it is defined as

$$\max J(t) = \frac{\left(\sum_{i\in\mathcal{V}(t)} \frac{P_i(t)}{w_i}\right)^2}{|\mathcal{V}(t)| \sum_{i\in\mathcal{V}(t)} (\frac{P_i(t)}{w_i})^2}. \quad (9)$$

Unlike the global utility (evaluated once per vehicle at exit), $J(t)$ is defined at all times and recomputed at every allocation step; we summarize a run by its mean $\bar{J}$.

Problem statement. At each event time t, the GSC allocates the budget A among the vehicles currently on the lane, i.e., it chooses $\{P_i(t)\}_{i\in\mathcal{V}(t)}$ subject to constraints (1) through (6), so as to jointly maximize the exit-level global utility (8) and the time-averaged fairness $\bar{J}$ derived from (9). Since these two objectives conflict, the problem is stated in the Pareto sense. Each allocation is decided from the current state alone: future arrivals being unknown, the schedulers optimize the instantaneous allocation, not the full trajectory.

IV. Scheduling Algorithms

Balancing utility and fairness in real time calls for heuristics that remain efficient under stochastic traffic. We focus on a tunable GU/fairness-greedy mixed scheduler (IV-E) and additionally consider a Weighted Max-Min Fair Allocation (IV-F), a classical fairness-oriented policy included as a natural candidate, together with two single-objective baselines for reference (IV-C, IV-D). All four rely on two common building blocks: the Mean Required Charging Power (IV-A) and a weighted water-filling routine (IV-B). Each scheduler uses only a constant number of sorts and water-filling passes, so every allocation runs in $O(n \log n)$, well within the real-time budget.

A. Mean Required Charging Power

We define the mean required charging power as the constant power that allows vehicle i to cover the remaining distance to its destination $(L + d_i^{\text{post}} - p_i(t))$ with sufficient energy.

Positions are measured from the lane entrance (origin), so the lane covers $[0, L]$ and d_i^{post} extends beyond the exit L. Since charging occurs only over sections, the remaining effective charging distance is the total overlap of $[p_i(t), L]$ with each section $[s_j^{\text{start}}, s_j^{\text{end}}]$:

$$L_i^{\text{rem}}(t) = \sum_{j=1}^{N_s} \max\left(0,\, s_j^{\text{end}} - \max\left(s_j^{\text{start}},\, p_i(t)\right)\right) \tag{10}$$

The mean required charging power then follows from the stored energy $y_i(t)B_i$ and the consumption c_i:

$$P_i^{\text{mreq}}(t) = \frac{(L + d_i^{\text{post}} - p_i(t))c_i - y_i(t)B_i v}{L_i^{\text{rem}}(t)} \tag{11}$$

If $P_i^{\text{mreq}}(t) \le 0$, the vehicle needs no further charging. By construction the last section ends at the lane exit, so $L_i^{\text{rem}}(t) > 0$ and $P_i^{\text{mreq}}(t)$ stays well defined; a vehicle with $P_i^{\text{mreq}}(t) > P_i^{\max}$, having the largest normalized demand, is served last and capped at $P_i^{\max}$.

B. Weighted Water-Filling

All four schedulers repeatedly solve the same elementary subproblem: share a fixed budget P among vehicles, each with a priority weight w_i and a cap κ_i, so as to equalize the weighted shares P_i/w_i as far as the caps allow. This is the classical *weighted water-filling* (progressive filling) allocation [20]: every share is raised at a common rate proportional to its weight,

$$P_i = \min(\kappa_i,\, \theta\, w_i)\,, \tag{12}$$

and the single water level θ is set so that the budget is exactly used up, $\sum_i \min(\kappa_i, \theta w_i) = P$. A vehicle whose cap κ_i is reached *saturates* while the remaining budget keeps filling the others; all non-saturated vehicles share the same normalized level $P_i/w_i = \theta$.

The level θ is found by sorting the vehicles once on κ_i/w_i and going through them in that order (Algorithm 1); the sort dominates, so a call costs $O(n \log n)$. We denote it WATERFILL$(\{(\kappa_i, w_i)\}, P)$. The same routine honors the per-section capacities a_j (constraint 2) at unchanged $O(n \log n)$ cost, omitted here for clarity.

Algorithm 1 Weighted water-filling, WATERFILL$(\{(\kappa_j, w_j)\},\, P)$

sort the items so that $\kappa_1/w_1 \le \cdots \le \kappa_n/w_n$
$C \leftarrow 0, \quad W \leftarrow \sum_j w_j$
for $k = 1$ **to** n **do**
  $\theta \leftarrow (P - C)\,/\,W$
  if $\theta \le \kappa_k/w_k$ **then**
    break ▷ level set; items $k, \ldots, n$ stay unsaturated
  end if
  $C \leftarrow C + \kappa_k, \quad W \leftarrow W - w_k$ ▷ item k saturates at κ_k
end for
return $\{\, P_j \leftarrow \min(\kappa_j,\, \theta\, w_j) : 1 \le j \le n \,\}$

C. Global-Utility Greedy Allocation Algorithm

The first baseline maximizes global utility. Vehicles are prioritized by their weight-normalized demand $P_i^{\text{mreq}}(t)/w_i$. Serving the smallest normalized demand first greedily maximizes the weighted number of satisfied users, ignoring fairness. Any power left once all of $\mathcal{V}(t)$ is served is redistributed by weighted water-filling (IV-B) among the still-active vehicles ($P_i^{\text{mreq}}(t) > 0$) up to $P_i^{\max}$, so none is sent to vehicles with no residual need. In the pseudocode, P_i denotes the i-th vehicle of S in sorted order, not an identifier.

Algorithm 2 Global-Utility Greedy Allocation

$S \leftarrow \{j \in \mathcal{V}(t) : P_j^{\text{mreq}}(t) > 0\}$, sorted by $P_j^{\text{mreq}}(t)/w_j$
$P_{rem} \leftarrow A, \quad P_j(t) \leftarrow 0,\ \forall j \in S, \quad i \leftarrow 0$
  ▷ Phase 1: greedy fill in ascending order of normalized demand P_j^{mreq}/w_j
while $P_{rem} > 0$ and $i < |S|$ **do**
  $alloc \leftarrow \min(P_i^{\text{mreq}}(t),\, P_i^{\max},\, P_{rem})$
  $P_i(t) \leftarrow alloc, \quad P_{rem} \leftarrow P_{rem} - alloc, \quad i \leftarrow i + 1$
end while
  ▷ Phase 2: water-fill the leftover over active vehicles, capped at $P_j^{\max}$
$R \leftarrow \{j \in S : P_j(t) < P_j^{\max}\}$
$\delta \leftarrow$ WATERFILL$(\{(P_j^{\max} - P_j(t),\, w_j) : j \in R\},\ P_{rem})$
$P_j(t) \leftarrow P_j(t) + \delta_j$ for all $j \in R$

It is worth mentioning that a globally optimal sequence would require non-causal knowledge of future states; this greedy heuristic is instead a strong, efficient utility-centric baseline.

D. Equal Allocation Algorithm

The second baseline maximizes Jain's index by equal allocation, equalizing $P_i(t)/w_i$ across all active vehicles subject to $P_i^{\max}$ and battery capacity. Ignoring individual needs, it is an upper bound on fairness and a benchmark for other policies. It is a direct weighted water-filling capped at the charging rate, WATERFILL$(\{(P_i^{\max}, w_i)\}, A)$, also $O(n \log n)$; no separate pseudocode is needed.

E. GU/Fairness-Greedy Mixed Algorithm

The GU/fairness-greedy mixed algorithm bridges the two objectives. It splits A: a fraction X follows the GU-greedy logic (IV-C), while $(1 - X)\, A$ (plus any power left unused by the greedy phase) is shared by the equal-allocation water-filling (IV-D) over all non-full vehicles. Varying X tunes the balance between utility and equity. Chaining the two building blocks, it needs no separate pseudocode and runs in $O(n \log n)$.

F. Weighted Max-Min Fair Allocation

The Weighted Max-Min Fair Allocation distributes the budget so as to equalize the normalized shares $P_i(t)/w_i$ across vehicles, each capped at the power it actually needs. It is obtained in two phases, both built on the weighted water-filling of Section IV-B.

The first, main phase water-fills A with every vehicle capped at its mean required power, but no higher than the hardware limit, $\kappa_i = \min\big(P_i^{\mathrm{mreq}}(t),\, P_i^{\max}\big)$, the shares growing together until each saturates (smallest κ_i/w_i first). This is exactly the weighted max-min fair allocation toward the mean required demands, computed directly by the water-filling in $O(n \log n)$. If power remains, a second phase water-fills it over the still-active vehicles, this time capped at $P_i^{\max}$, further equalizing $P_i(t)/w_i$ without disturbing the max-min structure.

Algorithm 3 Weighted Max-Min Fair Allocation

▷ Phase 1: weighted max-min toward the mean required demands
$S \leftarrow \{j \in \mathcal{V}(t) : P_j^{\mathrm{mreq}}(t) > 0\}$
$\{P_j(t)\} \leftarrow \textsc{WaterFill}(\{(\min(P_j^{\mathrm{mreq}}(t),\, P_j^{\max}),\, w_j) : j \in S\},\; A)$
▷ Phase 2: water-fill any residual power, capped at $P_j^{\max}$
$P_{rem} \leftarrow A - \sum_{j \in S} P_j(t), \quad R \leftarrow \{j \in S : P_j(t) < P_j^{\max}\}$
$\delta \leftarrow \textsc{WaterFill}(\{(P_j^{\max} - P_j(t),\, w_j) : j \in R\},\; P_{rem})$
$P_j(t) \leftarrow P_j(t) + \delta_j$ for all $j \in R$

This scheduler combines the spatially-aware urgency $P_i^{\mathrm{mreq}}(t)$ with weighted max-min fairness; like the others it runs in $O(n \log n)$ (Section V-D).

V. Performance Evaluation

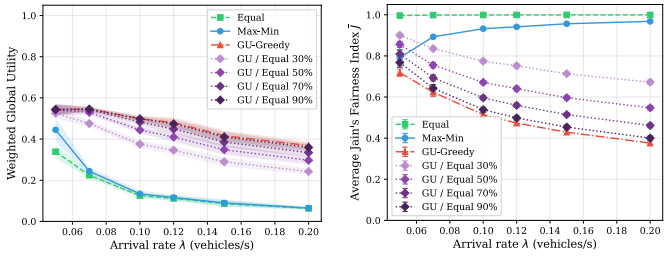
This section reports the simulation-based evaluation: the setup (V-A), the metrics (V-B), the results (V-C), and the end-to-end control-loop latency (V-D).

A. Simulation Setup

The scheduling algorithms are evaluated through an event-based simulation of the dynamic charging lane modeled in Section III. The simulation is implemented in Python 3.14 on a laptop (13th-Gen Intel Core i9-13950HX, 32 GiB DDR5, Ubuntu 26.04); the delays in Section V-D correspond to this configuration.

The simulation advances by irregular time steps to the next state change (the events listed in Section III-A); at each, the scheduler is triggered and its per-call wall-clock duration is measured with Python's `perf_counter`. To remove the start-up phase during which the road fills, metrics are recorded only once the simulated time exceeds $2L/v = 4$ h at $v = 50$ km/h (twice a vehicle's crossing time), beyond which the number of vehicles present has reached a stationary distribution and the statistics reflect steady-state operation.

The simulated infrastructure is a 100 km corridor with four 10 km charging sections at km $[0, 10]$, $[30, 40]$, $[60, 70]$, $[90, 100]$, of aggregate capacity $A = 5000$ kW with $a_j = 2000$ kW per section. After the warm-up transient fills the lane, the allocated power equals A at every step, so the system is strictly saturated (constraint 1). This partial electrification (four 10 km stretches out of 100 km) follows the eRoad concept of equipping major links rather than the whole network [1]. These are representative values placing the system in the saturated regime, not a specific deployment; the comparative



(a) Weighted global utility (b) Fairness $\bar{J}$

Fig. 2: Weighted global utility (a) and average weighted Jain's fairness index $\bar{J}$ (b) as a function of the arrival rate λ.

conclusions depend on the *degree of saturation* (varied through λ), not these magnitudes.

Passenger cars (90% of arrivals) follow the battery-capacity and consumption distributions of Weiss et al. [21] for the European EV fleet; as that reference omits charging power, each car's $P_i^{\max}$ is sampled uniformly in $[20, 50]$ kW, comparable to Qiu's $[40, 100]$ kW per-coil capacities [5]. Trucks (10%) are sampled uniformly from the eight commercial electric trucks of Liimatainen et al. [22], with $P_i^{\max}$ the DC charging power listed there (44 to 260 kW). Each priority weight is proportional to mean consumption, normalized to a car and rounded: $w_i = 5$ for medium-duty and rigid trucks, 12 for semitrailers, 1 for cars. The initial SOC is sampled uniformly in $[50\%, 80\%]$ and the post-lane distance d_i^{post} in $[150, 250]$ km.

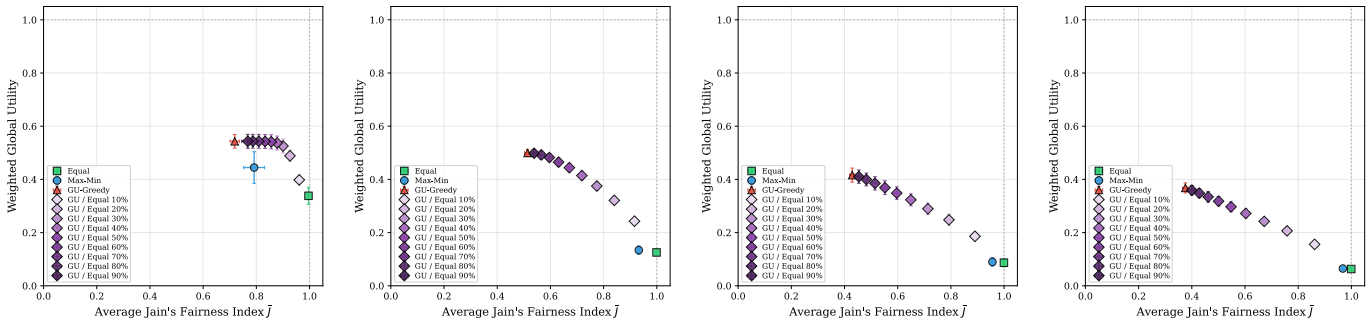
Two factors vary across scenarios: the algorithm and the arrival rate λ. Four algorithms are compared, the two baselines (equal and GU-greedy) and the two proposed heuristics (GU/fairness-greedy mixed and weighted max-min), with the mixed share X varied from 0.1 to 0.9 in steps of 0.1 and $\lambda \in \{0.05, 0.07, 0.10, 0.12, 0.15, 0.20\}$ veh/s, ranging from loaded to heavily saturated. Each (algorithm, λ) is replicated over $K = 10$ runs, each ending once 1000 vehicles have exited past the warm-up; seeds are shared across algorithms within a run, so they face identical arrivals. Metrics are aggregated over the 10 runs (Section V-B), giving 10 000 exits per pair.

B. Performance Metrics

Three metrics compare the algorithms, each computed per run and averaged over the $K = 10$ runs of a scenario. *Average weighted fairness* $\bar{J}$ averages $J(t)$ (9) over all post-warm-up allocation steps; its per-step standard deviation stays below 0.10, so the per-run mean accurately summarizes a run. *Weighted global utility* (8) is evaluated at each run's end over its 1000 post-warm-up exits as the satisfied-weight fraction. *Allocation delay* is each call's wall-clock duration, aggregated in Section V-D's boxplot over the 10 runs.

C. Results and Analysis

a) Global utility under load: Fig. 2a shows the weighted global utility versus λ. At every λ, GU-greedy is highest and equal allocation lowest (the latter ignoring residual energy needs), with the mixed family in between, increasing monotonically with X. The weighted max-min algorithm changes



(a) $\lambda = 0.05$ (b) $\lambda = 0.10$ (c) $\lambda = 0.15$ (d) $\lambda = 0.20$

Fig. 3: Tradeoff between fairness and utility at increasing load ($\lambda = 0.05, 0.10, 0.15, 0.20$ veh/s); the trend is monotonic across the remaining rates.

behavior across load: its utility (0.44 at $\lambda = 0.05$) lies between the baselines but drops to 0.065 at $\lambda = 0.20$, essentially matching equal allocation's 0.063.

b) Fairness under load: Fig. 2b reports $\bar{J}$ versus λ. Equal allocation reaches $\bar{J} \geq 0.99$ everywhere, as expected from its proportional share. GU-greedy is lowest, from 0.72 at $\lambda = 0.05$ to 0.38 at $\lambda = 0.20$, its sequential prioritization concentrating power on a subset. The mixed family is again intermediate, $\bar{J}$ decreasing with X. The weighted max-min shows the reverse trend to its utility curve: fairness grows from 0.79 at $\lambda = 0.05$ to 0.97 at $\lambda = 0.20$, confirming the load-dependent transition.

c) Tradeoff between fairness and utility: Fig. 3 plots every scenario in the $(\bar{J}, \text{utility})$ plane at various loads. As expected by construction, the mixed family traces a smooth continuous tradeoff curve joining the two baseline corners. The weighted max-min never lies on the useful frontier: it is strictly Pareto-dominated at every load by the mixed family (at $\lambda = 0.05$, $X \in \{0.3, \ldots, 0.7\}$ achieve higher fairness and utility), while at high loads it collapses onto the equal-allocation corner (at $\lambda = 0.20$ its fairness 0.97 nears equal allocation's, its utility advantage within run-to-run variation).

D. End-to-End Latency of the Control Loop

Reallocations close a real-time control loop (Section III-A). We assess their real-time viability by combining the measured execution time with V2I and wired delays from the literature.

a) Measured computation delay: Fig. 4 reports the per-step delay T_{alloc} per algorithm and λ. As expected, delays grow with λ as more vehicles are present. Equal allocation is cheapest (mean < 0.55 ms); weighted max-min and GU-greedy follow, their means rising from 0.27 and 0.42 ms at $\lambda = 0.05$ to 1.04 and 1.24 ms at $\lambda = 0.20$; the GU/fairness mixed family is the most demanding, chaining a greedy pass with a water-filling, yet stays below about 1.9 ms (at $\lambda = 0.20$). All four remain below 2 ms throughout, consistent with their $O(n \log n)$ complexity.

b) End-to-end latency budget: The computation delay is only one term of the control loop. Its end-to-end latency, from the instant a vehicle reports its state to the instant the granted power is applied, decomposes as

$$T_{\text{e2e}} = T_{\text{uplink}} + T_{\text{alloc}} + T_{\text{downlink}} \tag{13}$$

T_{uplink} and T_{downlink} each bundle a V2I air-interface and a wired backhaul term. For the air interface we use the field measurements of Charpentier *et al.* [11]: a mean OBU-to-RSU latency of 4.05 ms over ITS-G5 (802.11p), a conservative bound for 802.11bd, which roughly halves it [13]. The wired term, from the decomposition of Coll-Perales *et al.* [12], we take as a conservative 1 ms. Thus $T_{\text{uplink}} = T_{\text{downlink}} \approx 5$ ms, a communication overhead of about 10 ms.

The resulting end-to-end latency per algorithm at a mid-range load ($\lambda = 0.10$ veh/s) is reported in Table III: in all cases, the communication terms dominate the allocation term, with a maximum end-to-end delay of 11.0 ms.

TABLE III: Estimated end-to-end control-loop latency and vehicle displacement at $v = 50$ km/h ($\lambda = 0.10$ veh/s).

Algorithm	T_{alloc} (ms)	T_{e2e} (ms)	**Displacement**
Equal allocation	0.26	10.4	14.4 cm
GU-greedy	0.67	10.8	15.0 cm
GU/fairness mixed	0.91	11.0	15.3 cm
Weighted max-min	0.51	10.6	14.7 cm

c) Real-time compatibility: Converting these to physical distance, the per-iteration distance at $v = 50$ km/h ranges from 14.4 to 15.3 cm (Table III), four to five orders of

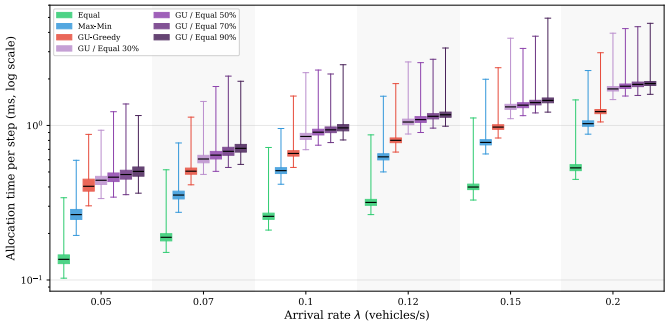


Fig. 4: Per-step allocation delay T_{alloc} distribution (logarithmic scale) as a function of the arrival rate λ.

magnitude below a section's 10 km, so the granted power is applied essentially where it was computed, well within the tens-of-milliseconds budget for V2X control [12]. Although this budget uses the mean latency, the narrow spread of T_{alloc} (Fig. 4) keeps even rare worst-case calls to a few centimeters, with no practical effect for this non-critical application.

VI. Conclusion

This paper reformulated real-time power scheduling on a saturated dynamic charging lane as a well-posed bi-objective problem (weight-normalized global utility vs. weighted Jain's fairness), and proposed two efficient heuristics against two single-objective baselines. In simulation, the mixed family traces a tunable Pareto frontier between fairness and utility at mean delays within about two milliseconds, Pareto-dominating the weighted max-min (which collapses onto equal allocation at high load). On the communication plane, by modeling the lane as a networked control system, we further showed that the control loop closes in about 11 ms over IEEE 802.11bd V2I links, well within the real-time budget.

Three limitations remain. First, the reported delays are upper bounds as they come from a reference Python implementation. Second, the focus is the saturated regime (constraint 1); unsaturated extensions would need to consider traffic prediction or demand elasticity. Third, the heuristics are intentionally simple, the formulation being the primary contribution; more advanced schedulers (short-horizon lookahead or learned policies [23]) would likely push the frontier further.

Two further directions concern deployment: the weights w_i represent a chosen fairness rule (here proportional to consumption, but could as well be equal access, willingness-to-pay, or remaining need) that the algorithms accept unchanged; and the self-declared vehicle profiles (B_i, c_i, d_i^{post}) could be misreported by self-interested drivers. These directions motivate mechanisms that control or reward honest reporting to avoid any potential abuse from the users.

Data Availability

The simulator, scheduling algorithms, and configuration files are available at https://github.com/gouthgouthgouth/powerAllocation. Per-run metrics are exported as CSV (one row per run-scenario tuple).

Acknowledgement

This work was supported by the eRoadMontBlanc project, funded by the French State under the France 2030 initiative, and by the European Union as part of the France Relance recovery plan, in collaboration with multiple public and private stakeholders: Alstom, ATMB, Greenmot and Pronergy.

References

[1] D. Connolly, "Economic viability of electric roads compared to oil and batteries for all forms of road transport," *Energy Strategy Rev.*, vol. 18, pp. 235–249, 2017.

[2] M. A. Khan, W. Burghout, O. Cats, E. Jenelius, and M. Cebecauer, "Charge-on-the-move solutions for future mobility: A review of current and future prospects," *Transp. Res. Interdiscip. Perspect.*, vol. 29, p. 101323, 2025.

[3] D. Danner and H. de Meer, "Quality of service and fairness for electric vehicle charging as a service," *Energy Informat.*, vol. 4, no. Suppl 3, p. 16, 2021.

[4] G. Fridgen, M. Thimmel, M. Weibelzahl, and L. Wolf, "Smarter charging: Power allocation accounting for travel time of electric vehicle drivers," *Transp. Res. D, Transp. Environ.*, vol. 97, p. 102916, 2021.

[5] C. Qiu, A. Sarker, and H. Shen, "Power distribution scheduling for electric vehicles in wireless power transfer systems," in *Proc. Int. Conf. Sensing, Commun., Netw. (SECON)*. IEEE, 2017, pp. 1–9.

[6] J. Li, B. Li, Q. Li, Z. Liu, and H. Su, "Co-optimized operation of coupled transportation and power networks considering mixed user equilibrium," in *Proc. Chinese Control Decis. Conf. (CCDC)*. IEEE, 2025, pp. 4544–4550.

[7] H. Liu, Y. Zou, Y. Chen, and J. Long, "Optimal locations and electricity prices for dynamic wireless charging links of electric vehicles for sustainable transportation," *Transp. Res. E, Logist. Transp. Rev.*, vol. 152, p. 102187, 2021.

[8] S. Zhang and J. J. Q. Yu, "Electric vehicle dynamic wireless charging system: Optimal placement and vehicle-to-grid scheduling," *IEEE Internet Things J.*, vol. 9, no. 8, pp. 6047–6057, 2022.

[9] Y. Wang, X. Li, M. Alhazmi, D. Xie, and X. Wang, "Optimizing wireless dynamic charging infrastructure for electric vehicles: A predictive and adaptive framework leveraging decision-dependent uncertainties and renewable energy utilization," *Int. J. Electr. Power Energy Syst.*, vol. 171, p. 111000, 2025.

[10] T. Bai, Y. Li, A. A. Malikopoulos, K. H. Johansson, and J. Mårtensson, "Distributed charging coordination for electric trucks under limited facilities and travel uncertainties," *IEEE Trans. Intell. Transp. Syst.*, 2025.

[11] V. Charpentier, N. Slamnik-Krijestorac, and J. Marquez-Barja, "Latency-aware C-ITS application for improving the road safety with CAM messages on the Smart Highway testbed," in *Proc. INFOCOM Wkshps.* IEEE, 2022, pp. 1–6.

[12] B. Coll-Perales *et al.*, "End-to-end V2X latency modeling and analysis in 5G networks," *IEEE Trans. Veh. Technol.*, vol. 72, no. 4, pp. 5094–5109, 2023.

[13] S. Xue, S. Gong, and X. Li, "A comparative study of IEEE 802.11bd and IEEE 802.11p on the data dissemination properties in dynamic traffic scenarios," *Appl. Sci.*, vol. 14, no. 5, p. 2099, 2024.

[14] G. Le Tat, I. Ahmim, A. Rachedi, A. Mebarki, M. Chachoua, and F. Vienne, "An adaptable digital twin architecture for smart and interoperable electrical roads," in *Proc. Consum. Commun. Netw. Conf. (CCNC)*. IEEE, 2026.

[15] I. Ahmim, F. Bouakkaz, A. Rachedi, and N. Ghoualmi-Zine, "LASeR: Lightweight and secure remote user authentication protocol for Internet of Drones," *J. Netw. Comput. Appl.*, vol. 242, p. 104275, 2025.

[16] I. Ahmim, N. Ghoualmi-Zine, F. Bouakkaz, and A. Rachedi, "2AS-DS: Anonymous authentication scheme based on physical unclonable function for drone swarms," in *Proc. Int. Conf. Embedded Distrib. Syst. (EDiS)*. IEEE, 2024, pp. 199–204.

[17] CAR 2 CAR Communication Consortium, "Next generation V2X – IEEE 802.11bd as fully backward compatible evolution of IEEE 802.11p," CAR 2 CAR Communication Consortium, White Paper C2CCC_WP_2098, V1.0, 2023.

[18] R. K. Jain, D.-M. W. Chiu, and W. R. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," Digital Equipment Corporation, Tech. Rep. DEC-TR-301, Sep. 1984.

[19] M. A. Beyazit, A. K. Erenoglu, and A. Tascikaraoglu, "Fairness and equity in electric vehicle charging with mobile charging stations," *IEEE Trans. Ind. Appl.*, vol. 61, no. 1, 2025.

[20] D. P. Bertsekas and R. G. Gallager, *Data Networks*, 2nd ed. Prentice-Hall, 1992.

[21] M. Weiss, T. Winbush, A. Newman, and E. Helmers, "Energy consumption of electric vehicles in Europe," *Sustainability*, vol. 16, no. 17, p. 7529, 2024.

[22] H. Liimatainen, O. van Vliet, and D. Aplyn, "The potential of electric trucks–an international commodity-level analysis," *Appl. Energy*, vol. 236, pp. 804–814, 2019.

[23] J. Zhou, Y. Xiang, S. Li, and W. Sun, "Learning to predictive scheduling for orderly charging of electric vehicles in renewable energy charging stations," *Renew. Energy*, vol. 256, p. 124332, 2026.