2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Path Planning for Four-Wheel Steering Forklifts in Constrained Spaces: A State-Embedded Hamiltonian Fast Marching Approach

Julien Pascal[1], Jean-Marie Mirebeau[2], Benoit Thuilot[1] and Paul Checchin[1]

***Abstract*— This paper proposes a global planner tailored to four-wheel steering (4WS) forklifts operating in constrained environments. It extends the Hamiltonian Fast Marching (HFM) to integrate 4WS non-holonomic constraints, multi-circle geometry, and a hybrid state system. This ensures mathematical guarantees of determinism, stability, and optimality, while fully exploiting 4WS maneuverability, properties that existing toolboxes and generic planners do not provide. Through qualitative and quantitative evaluations across various scenarios, we benchmark our method against leading motion planning algorithms, including RRT, RRT*, Informed-RRT*, and SST, using a standardized framework. Results demonstrate that our approach consistently outperforms existing solutions, particularly in maintaining minimum clearance distances and ensuring path predictability. The source code, along with modeling and implementation details necessary for replicating the results and benchmarking against state-of-the-art methods, is publicly available. `https://github.com/PascalJulien/State-HFM`**

I. Introduction

Industrial logistics and manufacturing environments demand the efficient movement of heavy goods across complex facilities (Fig. 1a). Material handling operations often involve navigating narrow aisles, densely occupied spaces, and performing precise tasks in confined areas. In this context, autonomous forklifts have become essential components of advanced logistics systems. Conventional autonomous vehicles typically use front-wheel or rear-wheel steering, which limits their maneuverability in tight spaces. The operational constraints of industrial environments require vehicles capable of executing complex paths while maintaining imposed alignment for lifting operations. Four-wheel steering (4WS) forklifts address these limitations by enabling independent control of both axles. This advanced steering configuration permits several unique motion capabilities that conventional vehicles cannot achieve: diagonal movement (crab steering) and reduced turning radius (counter-phase steering). This increased maneuverability is a challenge for traditional path planning algorithms, which were primarily designed for vehicles with simpler kinematic constraints (Fig. 1b). The industrial applications impose stringent requirements on path planning solutions. The algorithm must demonstrate deterministic behavior, producing identical paths under identical initial conditions to ensure operational predictability and repeatability. The solution must also maintain robust performance despite minor perturbations in environmental conditions, a common occurrence in dynamic industrial settings. The orientation requirement for load manipulation operations presents additional challenge. When approaching a pallet or container for lifting, the forklift must achieve target alignment with the load.

(a)

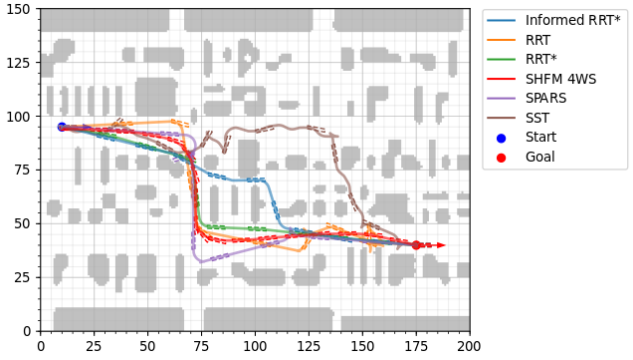

(b)

Fig. 1. (a) 4WS automated forklift in an industrial simulated environment, and (b) different paths obtained from state-of-the-art planning methods, considering 2WS forklifts with the case of a 4WS vehicle in red, for same initial (blue arrow with a dot) and target (red arrow with a dot) poses. The red path is (smooth, short, determinist).

This paper is organized as follows: Section II presents a comprehensive literature review and highlights the gaps our work aims to fill. Section III describes our proposed approach, including the detailed mathematical framework of our modified Hamiltonian Fast Marching (HFM) algorithm

This work is supported by the International Research Center "Innovation Transportation and Production Systems" of the I-SITE CAP 20−25. Financial support was also received from the Auvergne-Rhône-Alpes region through the ACCROBOT project (ACCostage haute précision ROBOTisé – Chantier Transitique du laboratoire FACTOLAB) as part of the Pack Ambition Recherche 2020. The funders had no role in the template design or decision to publish.

[1]Julien Pascal, Benoit Thuilot and Paul Checchin are with Institut Pascal, UMR 6602 CNRS, Université Clermont Auvergne `julien.pascal@doctorant.uca.fr`, `firstname.name@uca.fr`

[2]Jean-Marie Mirebeau is with Centre Borelli, ENS Paris-Saclay, UMR 9010 CNRS, University Paris-Saclay `jean-marie.mirebeau@ens-paris-saclay.fr`

and the state system design. Section IV outlines the experimental setup and evaluation framework, while Section V presents and discusses the results, offering insights into the algorithm's performance and broader implications for industrial automation.

II. Literature Review

Path planning constitutes a fundamental challenge for autonomous robotic systems, as evidenced by the recent study in [1], which references over 275 sources on the topic. This abundance of work reflects the critical importance of this NP-hard problem.Three main paradigms have emerged to solve this problem: graph-based methods such as Dijkstra [2], A* [3] or the Fast Marching Method (FMM) [4], which guarantee optimality but suffer from high computational costs; sampling-based approaches such as probabilistic roadmaps [5], which efficiently handle high-dimensional spaces at the cost of optimality; and gradient-based methods such as artificial potential fields [6], which offer computational efficiency but frequently encounter local minima issues. Recent developments such as the Traversability Hybrid A* algorithm [7] or A*/Bézier/Dynamic Window Approach combinations [8] attempt to improve path quality, while the Adaptive Dynamic Potential Field for Path Planning algorithm [9] seeks to overcome the inherent limitations of traditional gradient approaches. However, these generic approaches rely on simplifying assumptions that do not hold against the many constraints of industrial vehicles operating in highly constrained environments.

Non-holonomic vehicles, whose movements are subject to differential constraints, represent a particularly complex subset of the planning problem. The work of LaValle and Kuffner [10] with Rapidly exploring Random Trees (RRT) marked a significant advance specifically designed for these systems, while Mirebeau and Portegies [11] proposed an anisotropic FMM adapted to their constraints, henceforth referred to as the HFM method. Developments such as the attention-based sampling distribution for RRT [12] and the learning framework for sampling-based motion planning [13] have attempted to improve performance in autonomous driving scenarios with non-holonomic constraints. Furthermore, machine learning and deep learning techniques have been deployed both as components in the perception layer and as end-to-end driving solutions [14]. Guo et al. [15] developed an LSTM neural network architecture trained with data from a fuzzy control algorithm, while Kicki et al. [16] employed a gradient-based self-supervised learning algorithm, generating solutions 14 times faster than established methods such as RRT* and State Lattices [17]. Metaheuristic optimization algorithms demonstrate potential for path planning in autonomous driving. Recent surveys by Reda et al. [18] underscore the adaptability of these methods across various scenarios.

Although these approaches have demonstrated their effectiveness for conventional non-holonomic vehicles, they rely on a fundamental assumption: the existence of a relatively simple kinematic model with limited degrees of freedom. Moreover, the inherently non-deterministic nature of metaheuristic and machine learning approaches makes them unsuitable for industrial applications requiring consistent and repeatable path execution. While they could theoretically be applied to 4WS systems, current research has yet to explicitly address the unique challenges presented by these vehicles.

Autonomous forklifts equipped with four-wheel steering represent a particularly demanding case study that remains largely unexplored in the literature. Chen et al. [19] recently demonstrated that traditional path planning methods often neglect the kinematic richness of 4WS vehicles, leading to collision risks in complex environments. These vehicles possess a continuum of possible kinematic configurations, allowing them to dynamically adapt their movement to the geometric constraints of the environment. This kinematic flexibility, characterized by additional degrees of freedom, enables efficient navigation in restricted spaces inaccessible to conventional vehicles. Kokot et al. [20] introduced a generalized kinematic model for multi-wheeled AGVs. Yuan et al. [21] developed control laws for managing speed and steering angles in the context of tractor-trailer mobile robots. These contributions established theoretical foundations for variable orientation vehicles but focus primarily on control aspects rather than path planning.

Despite the significant advancements across multiple paradigms of path planning, a comprehensive examination of the literature reveals a critical gap in addressing the unique challenges posed by 4WS forklifts. These limitations can be categorized into the following key areas:

- *Deterministic guarantees*: Industrial environments require mathematically guaranteed repeatability, which stochastic methods cannot ensure.
- *Final pose alignment with the target orientation*: To conduct lifting operations, a consistent orientation of the forklift's forks is necessary when approaching.
- *Full usage of the agility*: Most existing methods lack the flexibility to leverage the enhanced maneuverability of 4WS robots when navigating through tight environmental constraints.
- *Predictable path generation*: To ease adoption by human operators, adjustable parameters ensure optimal paths remain predictable and intuitive.

To the best-of-our-knowledge, none of the existing approaches simultaneously addresses these requirements specific to the case of 4WS forklifts.

Faced with this clearly identified gap, our research proposes an extension of Mirebeau's work on FMM. Initially proposed by Tsitsiklis [22] for optimal control problems, and popularized shortly after by Sethian [4] for isotropic front propagation, FMM has revolutionized the approach to computing geodesic distances. It draws inspiration from Dijkstra's algorithm by simulating wavefront propagation from a source point and relies on first-order upwind discretization and node sorting based on arrival times, thus ensuring solution causality [23]. Mirebeau's work [24] marked a significant advancement by proposing an anisotropic FMM based on lattice basis reduction, an approach that naturally handles

anisotropic problems by locally adapting discretization to problem metric [25]. This method elegantly integrates non-holonomic constraints, making it particularly suitable for four-wheel steering vehicles.

This paper offers three significant advancements: (1) generalizing the model from 2WS to 4WS by integrating their specific non-holonomic constraints and extended kinematic space (Section III-C), (2) introducing a state system to further refine the final path (Section III-E), and (3) incorporating an efficient representation of the robot's shape (Section III-B). The presented approach maintains the mathematical guarantees of convergence, while fully leveraging the flexibility and kinematic adaptability of 4WS vehicles. Those properties are illustrated by different experiments and will be tested against state-of-the-art methods.

III. Proposed Approach

A. Kinematic model

The forklift state is represented by (x, y, θ) where (x, y) is the position of the rear axle center and θ is the orientation (see Fig. 2). Our goal is to find optimal paths respecting the vehicle's kinematic constraints. The forklift is modeled as a non-holonomic vehicle with kinematic constraints:

$$\begin{cases} \dot{x} = v\cos(\theta + \delta_R) \\ \dot{y} = v\sin(\theta + \delta_R) \\ \dot{\theta} = v\dfrac{\tan(\delta_F) - \tan(\delta_R)}{L} \end{cases} \tag{1}$$

where v represents the physical speed of motion, $\delta_R \in [-{\delta_R}^*; {\delta_R}^*]$ and $\delta_F \in [-{\delta_F}^*; {\delta_F}^*]$ are the rear and front steering angles respectively, whose maximum absolute value is denoted ${\delta_R}^*$ and ${\delta_F}^*$ respectively, and L is the wheelbase (distance between front and rear axles).

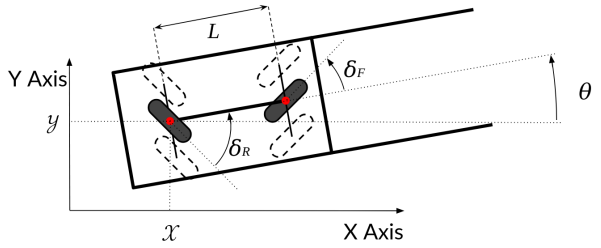

Fig. 2. Kinematic model of a 4WS forklift. L is the distance between the two axles, δ_F and δ_R are respectively the front and rear steering angle.

B. Representation of the robot footprint

The kinematic model describes the evolution of the rear axle center within $\mathbb{M} = \mathbb{R}^2 \times [0, 2\pi]$. A single enclosing circle (Fig. 3a) is inadequate for forklifts: the fork yields a 3:1 length-to-width ratio (Fig. 3b), wasting ~60% of the circle's area and blocking traversable passages. We therefore approximate the footprint with multiple overlapping circles (Fig. 3c), preserving $\mathcal{O}(1)$ collision queries with negligible overhead (Section V-C).

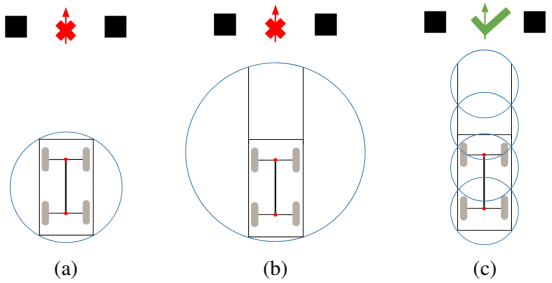

(a) (b) (c)

Fig. 3. In (a), a standard rectangular robot is roughly approximated by a circle. In (b), the presence of a fork demonstrates the inadequacy of this modeling approach. Finally, (c) presents an approximation using 4 circles.

C. Planning method

Our path planning problem requires finding minimum-cost paths while respecting the non-holonomic constraints of the forklift. This can be formulated as an anisotropic optimal control problem in the 3D configuration space $\mathbb{M} = \mathbb{R}^2 \times [0; 2\pi[$, with periodic boundary conditions in the angular dimension. Unlike the classical FMM which assumes isotropic motion (equal cost in all directions), we must adapt the method to accommodate the directional constraints of the forklift. Therefore, we need to account for the third angular dimension, and only allow kinematically admissible velocities in the three-dimensional configuration space.

We adopt a generative approach, and fix a family $(\dot{p}_i(q))_{1\le i\le I} \in \mathbb{R}^3 = (\dot{x}_i, \dot{y}_i, \dot{\theta}_i)$ of kinematically admissible velocities of the vehicle in configuration space at each point $q = (x, y, \theta) \in \mathbb{M}$. For a 4WS vehicle (1), the simplest choice is to use a unit physical speed $v_{4WS} = 1$ and consider the velocities $(\dot{p}_i^{4WS}(q))_{1\le i\le 4}$ associated with the four configurations of maximum steering of the front $\delta_F = \pm{\delta_F}^*$ and rear wheels $\delta_R = \pm{\delta_R}^*$:

$$\begin{bmatrix} \cos(\theta + {\delta_R}^*) \\ \sin(\theta + {\delta_R}^*) \\ \frac{\tan {\delta_F}^* - \tan {\delta_R}^*}{L} \end{bmatrix} \begin{bmatrix} \cos(\theta + {\delta_R}^*) \\ \sin(\theta + {\delta_R}^*) \\ \frac{-\tan {\delta_F}^* - \tan {\delta_R}^*}{L} \end{bmatrix} \begin{bmatrix} \cos(\theta - {\delta_R}^*) \\ \sin(\theta - {\delta_R}^*) \\ \frac{\tan {\delta_F}^* + \tan {\delta_R}^*}{L} \end{bmatrix} \begin{bmatrix} \cos(\theta - {\delta_R}^*) \\ \sin(\theta - {\delta_R}^*) \\ \frac{-\tan {\delta_F}^* + \tan {\delta_R}^*}{L} \end{bmatrix}$$

These directions are only a basis for the PDE update: the value function smooths them into continuous gradients, enabling smooth path extraction.

In Section III-E, we consider other modes of operations of the forklift, referred to as states. In the forward navigation (2WS) state, the rear wheels are blocked ($\theta_R = 0$), but the vehicle speed increases, especially when in straight line: $v_{\text{FWD}} > v_{2WS} > v_{4WS} = 1$. This mode is characterized by three velocities $(\dot{p}_i^{2WS}(q))_{1\le i\le 3}$ defined as:

$$v_{2WS} \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \\ \frac{\tan {\delta_F}^*}{L} \end{bmatrix}, \quad v_{2WS} \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \\ \frac{-\tan {\delta_F}^*}{L} \end{bmatrix}, \quad v_{\text{FWD}} \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \\ 0 \end{bmatrix}.$$

We also consider backward 4WS and 2WS maneuvering states, where the vehicle moves in reverse, which are defined by opposite velocities $(-\dot{p}_i^{4WS}(q))_{1\le i\le 4}$ and $(-\dot{p}_i^{2WS}(q))_{1\le i\le 3}$ to the above. In the rest of this section and until III-E, we assume that the state (operating mode) of the vehicle is fixed, and we denote by $(\dot{p}_i(q))_{1\le i\le I}$ the chosen kinematically admissible velocities.

A convenient way to visualize a non-holonomic vehicle's maneuvering capability is to draw Tissot's indicatrix (Fig. 4)

$$B(q) := \mathrm{Hull}\big(\{\dot{p}_i(q)\ ;\ 1 \le i \le I\} \cup \{0\}\big).$$

Let us fix a domain $\Omega \subset \mathbb{M} = \mathbb{R}^2 \times [0; 2\pi[$, collecting all configurations where the vehicle is authorized to go and does not intersect walls. Consider also an environmental cost factor $c : \Omega \to]0; \infty[$ which penalizes e.g. closeness to the obstacles. We say that a trajectory $q : [0; T] \to \Omega$ is admissible if it is Lipschitz continuous and

$$c(q(t))\,\dot{q}(t) \in B(q(t))$$

for all $t \in [0; T]$. Consider the minimal time optimal control problem, of moving from a fixed starting configuration $q_0 \in \Omega$, to an arbitrary endpoint $q_1 \in \Omega$. The *value function* reads

$$U(q_1) := \inf\{T \ge 0 \mid \exists q : [0; T] \to \Omega,\ q(0) = q_0,\ q(T) = q_1\},$$

where the path q is implicitly assumed to be admissible. Under mild assumptions [26], the value function $U : \Omega \to [0; \infty]$ is the viscosity solution to the anisotropic eikonal equation (or static Hamilton-Jacobi-Bellman equation):

$$\forall q \in \Omega \setminus \{q_0\}, \quad \max_{1 \le i \le I} \big\langle \nabla U(q), \dot{p}_i(q) \big\rangle = c(q), \qquad (2)$$

subject to the point source condition $U(q_0) = 0$, and to the outflow boundary conditions $U(q) = +\infty$ for all $q \in \partial\Omega$. The value function $U : \Omega \to [0, \infty]$, also known as the navigation function, constitutes the cornerstone of the path planning formulation. It defines a scalar field across the configuration space (Fig. 5) whose gradient directions indicate optimal paths that simultaneously respect both the forklift's non-holonomic constraints and environmental boundaries. Determining U is essential for path extraction, which we accomplish by numerically solving equation (2) using the FMM, providing an efficient computational approach with mathematical guarantees of convergence [25].

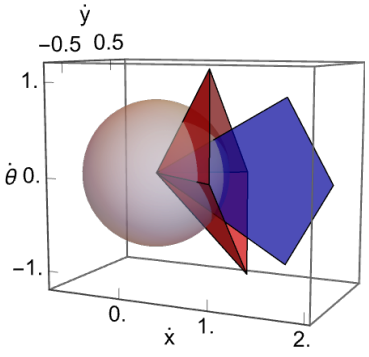

Fig. 4. Comparative view of Tissot indicatrices for three vehicle types: (gray) Holonomic vehicle can move in any direction (spherical indicatrix), (blue) Car-like vehicle with 2-wheel steering has more restricted movement (narrower shape), (red) Forklift with 4-wheel steering has enhanced maneuverability (diamond-like shape).

D. Numerical solution

The numerical computation of the value function requires the discretization of the domain Ω, and the approximation of the partial differential equation (2) using finite differences. The configuration space is discretized by sampling points regularly in a grid-pattern X on each dimension, whose gridscale is denoted $h > 0$. Those sampled points are noted $\mathbf{q} \in X$, and we assume that $q_0 \in X$. We discretize the PDE (2) using adaptive finite differences, following the HFM framework [11, 25]: for all $\mathbf{q} \in X \setminus \{q_0\}$

$$\max_{1 \le i \le I} \sum_{1 \le j \le J} \rho_{ij}(\mathbf{q}) \max\Big\{0, \frac{\mathbf{U}(\mathbf{q}) - \mathbf{U}(\mathbf{q} - h\dot{e}_{ij})}{h}\Big\}^2 = c(\mathbf{q})^2 \quad (3)$$

This formulation requires special attention to the weights $\rho_{ij}(\mathbf{q}) \ge 0$ and the offsets $\dot{e}_{ij} \in \mathbb{Z}^3$, whose design is non-trivial, see [25] for details. Their construction is necessary since the point $\mathbf{q} - h\dot{p}_i$ has no guarantee to be on the discrete grid. They serve to reconstruct the displacement $-h\dot{p}_i$ from $\mathbf{q}$ using points on the grid. This reconstruction technique is the base to anisotropic FMM. The integer I, which also appears in the original PDE (2), denotes the number of kinematically admissible velocities considered, and the integer $J = 6$ when the offsets depend on the current point, i.e. $\dot{e}_{ij} = \dot{e}_{ij}(\mathbf{q})$, see [11]. The contribution to (3) associated to a pair (i, j) of indices is omitted if $\mathbf{q} - h\dot{e}_{ij}$ lies outside the computational domain, or if the line segment $[\mathbf{q}, \mathbf{q} - h\dot{e}_{ij}]$ intersects an obstacle (*visibility condition*).

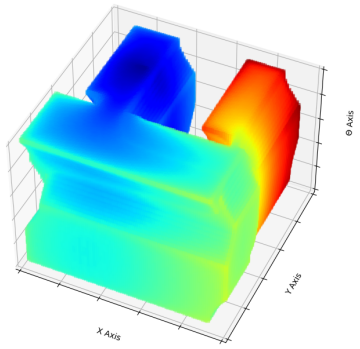

Fig. 5. Display of the navigation function U (here, related to the environment of Fig. 9). The color transitions from blue to red as we move further from the initial configuration. The color gradient indicates the path to follow to reach any configuration starting from the initial one.

This scheme obeys essential monotonicity and causality properties, allowing solution via FMM, see Alg. 1. The algorithm processes grid points in order of increasing values, similar to Dijkstra's algorithm for shortest paths on graphs. The difference lies in the *update* at $\mathbf{q} \in X \setminus \{q_0\}$, which consists in solving (3) w.r.t. the variable $U(\mathbf{q})$, while regarding the neighbor values $[U(\mathbf{q} - h\dot{e}_{ij})]_{1 \le j \le J}^{1 \le i \le I}$ as fixed, see [27, Alg. 3]. Alternatively to the sequential FMM described in Alg. 1, a parallel GPU accelerated solver is also available [27].

Once the value function $\mathbf{U}$ is computed, optimal paths can be extracted by following the gradient flow of U from any starting configuration to the seed point q_0. The resulting paths

properly respect the vehicle's non-holonomic constraints, but may still lack predictability for human operators.

Algorithm 1: FMM planner for 4WS forklift.

Input: X discrete domain; $c(\mathbf{q}) > 0$ cost function; $\dot{p}_i(\mathbf{q}) \in \mathbb{R}^3$ kinematically admissible vectors, $1 \le i \le I$

Output: $\mathbf{U}$ solution to (3)

Pre-compute weights $\rho_{ij}(\mathbf{q})$, offsets $\dot{e}_{ij}$, and reversed stencils $V(\mathbf{q}) := \{\mathbf{p} \in X; \exists i, j, \mathbf{q} = \mathbf{p} - h\dot{e}_{ij}\}$ for all $\mathbf{q} \in X$;
Initialize $\mathbf{U}(q_0) = 0$, $\mathbf{U} = +\infty$ elsewhere; tag q_0 as Front, others as Far;
while $\exists$ *Front points* **do**
  $\mathbf{q} \leftarrow \min_{\mathbf{p} \in \text{Front}} \mathbf{U}(\mathbf{p})$; tag $\mathbf{q}$ as Frozen;
  for $\mathbf{p} \in V[\mathbf{q}]$ **do**
    If $\mathbf{p}$ is Far, tag as Front;
    Update $\mathbf{U}[\mathbf{p}]$ using (3);

E. State system

To produce more intuitive and human-like paths, we integrate a state system $\mathbb{S} = \{0, \dots, S-1\}$ representing different operational modes. This new model falls in the framework of the optimal control of hybrid switched systems, see the survey [28] and recent works [29]. To our knowledge, this is the first time the anisotropic fast marching method is combined with the switching system framework, so as to compute the optimal trajectories of a non-holonomic vehicle with different operating modes. Unlike a heuristic bias, this is simply a higher-dimensional anisotropic eikonal problem: solving in (x, y, θ, s) ensures that global optimality is preserved.

Using only 4WS controls everywhere would produce paths that interleave crab motions unpredictably, even in open corridors as shown in Fig. 6(a): without the state system, the planner uses crab motion for over 90% of the trajectory, including counter-intuitive lateral moves in open areas. The state system resolves this, navigation states favor straight motion ($v_{\text{FWD}} > v_{2WS} > v_{4WS}$), so 4WS activates only when geometrically beneficial. Each state s (Fig. 7) has a specific set of controls $\dot{p}_{is}(q)$ (directional vectors), see the discussion in Section III-C. A transition cost $T(s, s')$ represents the discrete operational effort required to switch between states and are thus constant scalars (not functions of velocity or steering angle) forming a symmetric matrix.

The governing PDE is therefore: for each $(q, s) \in \Omega \times \mathbb{S} \setminus \{(q_0, s_0)\}$

$$\max\left\{\min_{1 \le i \le I} \frac{\langle \nabla U, \dot{e}_{is}(q)\rangle}{c(q)}, \min_{s' \ne s} \frac{U(q,s) - U(q,s')}{T(s,s')}\right\} = 1 \quad (4)$$

together with the point source $U(q_0, s_0) = 0$, and outflow boundary conditions $U = +\infty$ on $\partial\Omega \times \mathbb{S}$. This is not a heuristic, the embedding in (4) preserves global optimality.

The spatial part of (4) is discretized using adaptive finite differences similar to (3), whereas the state part is already in suitable finite difference form. This new numerical scheme leads to the following discrete update rule:

$$U(q,s) = \min\{U_{spatial}, U_{state}\} \quad (5)$$

with $U_{spatial}$ representing the optimal value obtained through spatial motion within state s, computed similarly to the spatial-only case (3), and $U_{state} = \min_{s' \ne s}\{U(q,s') + T(s',s)\}$, the optimal value achieved by switching states at q. The state system produces more predictable paths that naturally alternate between direct movements and complex maneuvers, similar to those created by human operators, while maintaining mathematical optimality. Fig. 6(b) illustrates the benefits: 4WS maneuvering is reserved for tight spaces, yielding human-like paths.

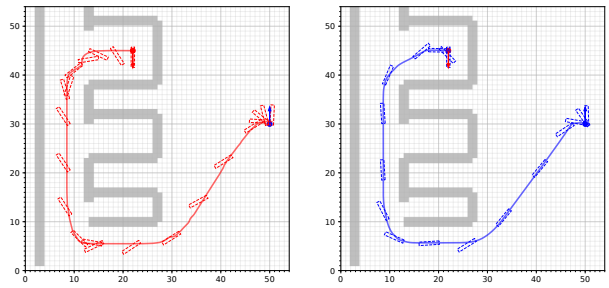

(a) (b)

Fig. 6. Navigation (a) without and (b) with the state system. Both paths are collision-free, but (a) uses crab motion throughout the generated path, while (b) reserves 4WS maneuvering for tight areas only.

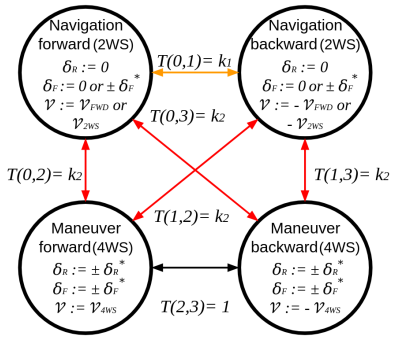

Fig. 7. State system for our forklift. The vectors $\dot{p}_{is}(q)$ of each state are the ones generated using the kinematic model with the specified parameters (δ_R, δ_F, v) in each state.

IV. Experimental Setup

The method takes as input a binary occupancy map, vehicle geometry, start/goal configurations, and the state system parameters. The cost function $c(q) = 1/(d(q) + \varepsilon)$ penalizes proximity to obstacles, with a saturation $c_{D_s}(q) = \max(c(q), 1/D_s)$ to mitigate excessively large deviations in open spaces. The footprint uses 12 circles and the grid resolution is $(N_x, N_y, N_\theta) = (200, 200, 100)$. Starting and ending points can be specified as coordinates (x, y), configurations (x, y, θ) or state-specific configurations (x, y, θ, s).

The state system comprises two navigation states (forward/backward, 2WS) and two maneuver states (forward/backward, 4WS) (Fig. 7). We choose to have a cost symmetry for the transition between state hence $T(s,s') = T(s',s)$. The transition between the forward and backward navigation state has a cost k_1 inferior to the cost to switch

between any of the navigation state to any of the maneuvering sate k_2: switching direction within navigation mode is cheap, while entering maneuver mode incurs a higher penalty, reserving 4WS for geometrically necessary situations. Finally, the cost to alternate motion directions once in a maneuver state is unitary. Navigation states use three control vectors (turn left, turn right, straight) with higher norms for straight motion to encourage straight-line motion when possible. Moreover, we have defined a backward velocity v_{BWD} inferior to forward velocity v_{FWD} to penalize backward motion. Maneuver states use four vectors at extreme δ_R, δ_F values with lower norms, discouraging their use unless required.

V. Results

To evaluate our approach, we conduct comprehensive testing in various environments (Fig. 8 & 9) and procedurally generated scenarios. Quantitative comparisons against state-of-the-art path planning methods are provided using the Bench-MR tool [30]. This framework enables systematic evaluation against prominent algorithms including RRT [31], RRT* [32], Informed-RRT* [33], and SST [34] which form the core of common toolboxes such as ROS NavStack [35]. Performance is assessed using standard Bench-MR metrics [30]: Path Length, Normalized Curvature, Maximum Curvature, and Minimum Clearing. We additionally introduce two custom metrics, namely Angular Deviation (AD) and Euclidean Deviation (ED), measuring approach stability toward the target pose $q^* = (x^*, y^*, \theta^*)$. Let $\bar{\ell}$ be a small displacement, and q^{*1}, q^{*2} and q^{*3} three configurations at $\bar{\ell}$-distance from each other starting from q^*, we define:

$$\text{AD} = \frac{\Delta\theta(q^*, q^{*1}) + \Delta\theta(q^{*1}, q^{*2}) + \Delta\theta(q^{*2}, q^{*3})}{3\,\bar{\ell}}, \quad (6)$$

$$\text{ED} = \frac{\|q^*_{xy} - q^{*1}_{xy}\| + \|q^{*1}_{xy} - q^{*2}_{xy}\| + \|q^{*2}_{xy} - q^{*3}_{xy}\|}{3\,\bar{\ell}}, \quad (7)$$

where $\Delta\theta(q^{*i}, q^{*j})$ denotes the angular difference modulo 2π. These metrics quantify convergence rate in orientation and position during the final approach; lower values indicate smoother arrival.

A. Advantage of the 4WS

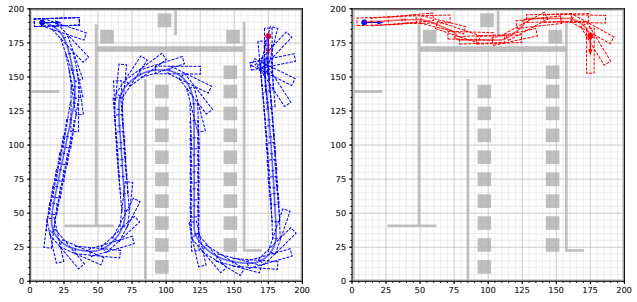

(a) (b)

Fig. 8. Comparison of path planning outcomes in a structured environment designed to highlight 4WS advantages. (a) The 2WS forklift must take the longer lower route due to its more limited steering capabilities, while (b) the 4WS configuration successfully navigates the narrow upper passage.

Fig. 8 compares 4WS and 2WS configurations. The 2WS case (Fig. 8a) applies the original HFM of Mirebeau and Portegies [11] with our state system restricted to front-wheel steering, serving as a direct baseline. The 4WS forklift (Fig. 8b) navigates the narrow upper passage by switching between navigation and maneuver states, while the 2WS configuration must take the longer lower route: its Tissot indicatrix (Fig. 4) lacks the lateral motions required. This illustrates how the enriched kinematic space unlocks paths inaccessible to standard HFM.

B. Adjustable parameters

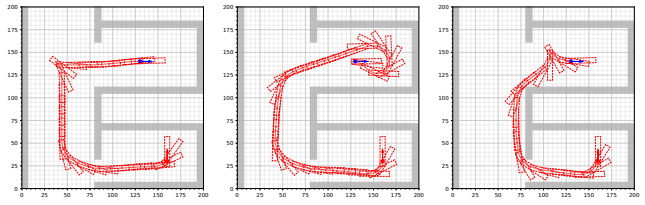

(a) (b) (c)

Fig. 9. Sensitivity to key parameters: (a) halving backward speed concentrates maneuvers at unavoidable locations, (b) large transition costs k_2 suppress maneuvering for wide arcs, (c) removing obstacle costs ($D_s \to \infty$) prioritizes geometry over clearance.

Strategic configuration of control parameters enables predictable paths tailored to operational requirements. As illustrated in Fig. 9, adjusting the backward velocity ratio $v_{\text{BWD}}/v_{\text{FWD}}$, state transition costs k_1, k_2, or obstacle saturation distance D_s produces qualitatively different but consistently smooth paths. This tunability steers the planner toward operator-preferred behavior without compromising optimality, a critical factor for trust in automated forklifts. Small parameter perturbations yield proportionally small path deviations, confirming stability.

C. Comparison with state-of-the-art methods

We evaluate with extensive comparisons against state-of-the-art path planning algorithms [30], which we extended to support 4WS kinematics by incorporating reduced turning-radius primitives and additional crab-motion primitives into their steering functions. This ensures that baselines can exploit the same lateral and tight-turning capabilities as our method, making the comparison fair. Quantitative evaluation encompasses both the same test environment as the precedent paragraph and procedurally generated scenarios (Fig. 12), with consistent performance metrics applied across all methods (Fig. 11 & 13).

A first distinction is determinism: our method always produces the same path for identical inputs (Fig. 10), while sampling-based planners show run-to-run variability. This consistency is critical in industrial deployments. This determinism is further quantified in Fig. 11, where the metrics for our approach shows zero deviation between runs.

Next, simply equipping existing planners with 4WS steering does not yield systematic gains: they lack mechanisms to balance navigation and maneuvering modes, often producing unnatural paths. State HFM's state system activates 4WS only when beneficial, preserving predictable straight motion otherwise.

Analysis of the test environment metrics in Fig. 11 reveals that our State HFM method outperforms sampling-based

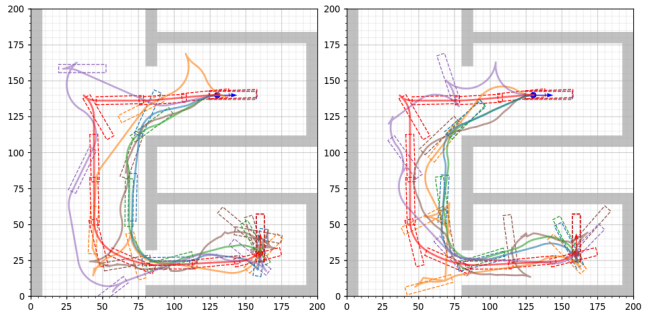

Fig. 10. Different executions of path planning algorithms in test environment. Our State HFM 4WS (red) methods consistently produce the same path for identical inputs, unlike sampling-based methods which exhibit variability. The color code of other methods is indicated in Fig. 12.

approaches across most performance indicators. While path length remains competitive (though longer than some alternatives), our approach demonstrates superior performances. Particularly noteworthy are the Euclidean Deviation Metrics and Angular Deviation Metrics, which show substantially better results for our method. The lower values achieved by our approach indicates smoother, more stable final positioning. Likewise, the minimum clearance distance is large, which is desirable. The maximum curvature metric is high because State HFM aims to minimize maneuvering segments; however, the normalized curvature remains competitive with the other algorithms, indicating that the generated paths are overall smooth.

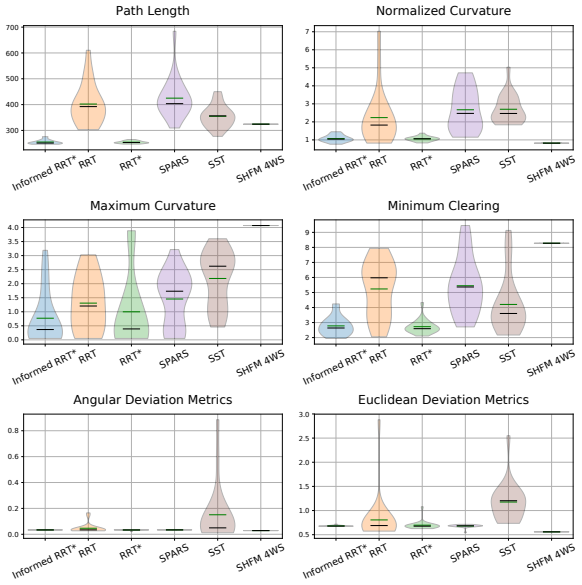

Fig. 11. Statistics of the different metrics on 20 executions on the environment of Fig. 10. The green bar indicates the mean value and black the median one. We see no deviation on State HFM since it is entirely deterministic. Running the solver on the same environment with the same extreme points give the same results.

To ensure robust evaluation across diverse scenarios, we extend our evaluation to procedurally generated environments. Fig. 12 presents a collection of corridor-like configurations with varying complexity. The paths produced by different planners are overlaid, with our State HFM 4WS (red) methods clearly visible. Hence, our approach maintains consistent behavior across different environmental configurations, while sampling-based methods exhibit higher variability.

The comprehensive metrics across 20 executions in these procedurally generated environments, shown in Fig. 13, provide compelling evidence of our method's superiority. State HFM consistently outperform other state-of-the-art methods in most metrics. Remarkable improvements in minimum clearing distances from obstacles are obtained, up to 200% in certain scenarios. The normalized curvature and maximum curvature metrics reveal smoother paths compared to sampling-based alternatives.

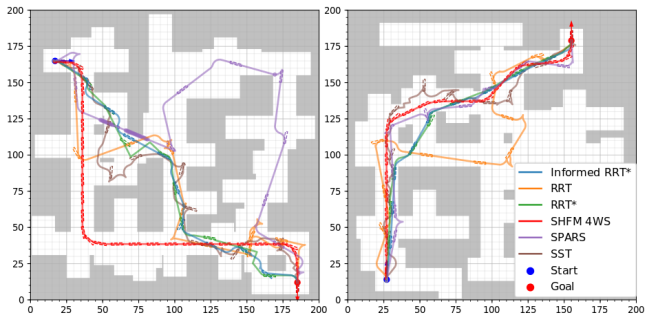

Fig. 12. Two of procedurally generated corridor-like environments used in our evaluation. Paths generated by different planners are shown, with our and State HFM 4WS (red) methods clearly visible. Each subplot represents a unique environment configuration generated with different random seeds.

Qualitative analysis of the generated paths (Fig. 12) further validates our approach's advantages. The visual representations clearly illustrate how our method produces cleaner, more intuitive paths which maintain appropriate safety margins. Particularly noteworthy is the consistent performance across different environmental configurations, while sampling-based methods exhibit higher variability.

In terms of computational cost, the overhead of multi-circle collision checking grows linearly with the number of circles and is negligible in practice. State HFM solves the full grid $(200\times200\times100\times4)$ in $\sim$3.6 s on an Intel i7-11850H (2.50 GHz, integrated GPU). This is a one-time cost, after which each query only requires gradient descent for path extraction (<0.01 s). By contrast, sampling-based approaches incur per-query runtimes from 2.38 s (RRT*) to 3.88 s (SST).

VI. Conclusion

In this paper, we introduce State Hamiltonian Fast Marching (HFM), a 4WS-specific global planner with deterministic and mathematically guaranteed trajectories, unlike existing extensions of generic methods. Built on the robust mathematical framework of Hamiltonian Fast Marching, our results confirm advantages in stability, clearance, and predictability, making State HFM a strong foundation for industrial forklifts. All results presented in this paper are obtained using publicly available code, ensuring reproducibility.

Several promising avenues for future research stem from this work. The primary focus is physical implementation on real forklift platforms to validate performance under real-world conditions (surface textures, sensor noise, actuator dynamics). Extending the approach to dynamic obstacles through time-dependent cost terms would broaden applicability. The framework could also generalize to other vehicles with complex kinematics, such as articulated or specialized industrial equipment.

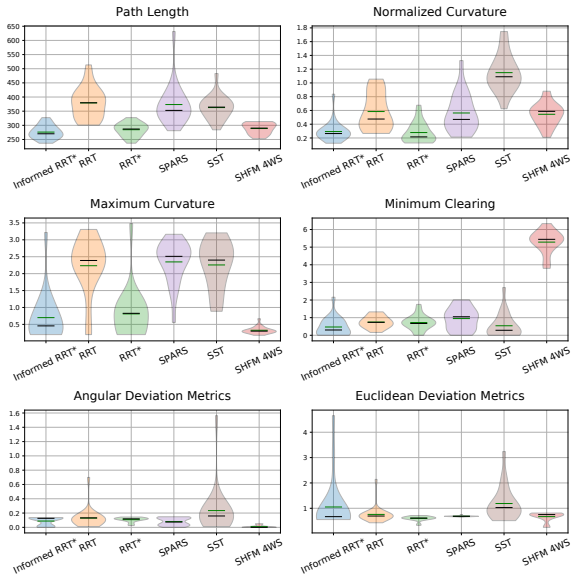

Fig. 13. Statistics of the different metrics on 20 executions on different procedurally generated environments. The green bar indicates the mean value and black the median one. State HFM outperform other state-of-the-art methods. The only exception being the path length, but we remain competitive. The trade-off between slightly longer paths and significantly better safety metrics represents an appropriate balance for industrial applications.

REFERENCES

[1] M. Reda *et al.*, "Path planning algorithms in the autonomous driving system: A comprehensive review," *Robot. and Auton. Syst.*, vol. 174, p. 104630, 2024.

[2] D. Delling *et al.*, "PHAST: Hardware-Accelerated Shortest Path Trees," in *IEEE Int. Parallel & Distrib. Processing Symp.*, 2011, pp. 921–931.

[3] P. E. Hart *et al.*, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," *IEEE Trans. on Syst. Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[4] J. A. Sethian, "A fast marching level set method for monotonically advancing fronts," *Proc. of the Natl. Acad. of Sci.*, vol. 93, no. 4, pp. 1591–1595, 1996.

[5] L. E. Kavraki *et al.*, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Trans. on Robot. and Autom.*, vol. 12, no. 4, pp. 566–580, 1996.

[6] O. Khatib, "The Potential Field Approach And Operational Space Formulation In Robot Control," in *Adaptive and Learning Syst.: Theory and Applications*, K. S. Narendra, Ed. Boston, MA: Springer US, 1986, pp. 367–377.

[7] M. Thoresen *et al.*, "Path Planning for UGVs Based on Traversability Hybrid A*," *IEEE Robot. and Autom. Lett.*, vol. 6, no. 2, pp. 1216–1223, 2021.

[8] Y. Li *et al.*, "A Mobile Robot Path Planning Algorithm Based on Improved A* Algorithm and Dynamic Window Approach," *IEEE Access*, vol. 10, pp. 57 736–57 747, 2022.

[9] B. Lu *et al.*, "Adaptive Potential Field-Based Path Planning for Complex Autonomous Driving Scenarios," *IEEE Access*, vol. 8, pp. 225 294–225 305, 2020.

[10] S. M. LaValle and J. J. Kuffner, "Randomized Kinodynamic Planning," *The Int. J. of Robot. Research*, vol. 20, no. 5, pp. 378–400, 2001.

[11] J.-M. Mirebeau and J. Portegies, "Hamiltonian Fast Marching: A numerical solver for anisotropic and non-holonomic eikonal PDEs," *Image processing on line*, vol. 9, pp. 47–93, 2019.

[12] J. Rong *et al.*, "Attention-based Sampling Distribution for Motion Planning in Autonomous Driving," in *39th Chinese Control Conf.*, 2020, pp. 5671–5676.

[13] Y. Zhang *et al.*, "A Novel Learning Framework for Sampling-Based Motion Planning in Autonomous Driving," *Proc. of the AAAI Conf. on Artificial Intelligence*, vol. 34, no. 01, pp. 1202–1209, 4 2020.

[14] X. Xiao *et al.*, "Motion planning and control for mobile robot navigation using machine learning: a survey," *Auton. Robots*, vol. 46, no. 5, pp. 569–597, 2022.

[15] N. Guo *et al.*, "Local Path Planning of Mobile Robot Based on Long Short-Term Memory Neural Network," *Automatic Control and Computer Sci.*, vol. 55, no. 1, pp. 53–65, 2021.

[16] P. Kicki *et al.*, "A Self-Supervised Learning Approach to Rapid Path Planning for Car-Like Vehicles Maneuvering in Urban Environment," 2020. [Online]. Available: https://arxiv.org/abs/2003.00946

[17] M. Pivtoraiko *et al.*, "Differentially constrained mobile robot motion planning in state lattices," *J. of Field Robot.*, vol. 26, no. 3, pp. 308–333, 2009.

[18] M. Reda *et al.*, "A discrete variant of cuckoo search algorithm to solve the Travelling Salesman Problem and path planning for autonomous trolley inside warehouse," *Know.-Based Syst.*, vol. 252, no. C, Sep. 2022.

[19] Y.-L. Chen *et al.*, "A path planner enabling second-order continuity across multiple motion modes of the generalized bicycle model," *Int. J. of Adv. Robotic Syst.*, vol. 21, no. 5, p. 17298806241277195, 2024.

[20] M. Kokot *et al.*, "Path continuity for multi-wheeled AGVs," *IEEE Robot. and Autom. Lett.*, vol. 6, no. 4, pp. 7437–7444, 2021.

[21] J. Yuan *et al.*, "Trajectory generation and tracking control for double-steering tractor–trailer mobile robots with on-axle hitching," *IEEE Trans. on Industrial Electronics*, vol. 62, no. 12, pp. 7665–7677, 2015.

[22] J. N. Tsitsiklis, "Efficient algorithms for globally optimal trajectories," *IEEE Trans. on Autom. Control*, vol. 40, no. 9, pp. 1528–1538, 1995.

[23] J. A. Sethian and A. Vladimirsky, "Ordered Upwind Methods for Static Hamilton–Jacobi Equations: Theory and Algorithms," *SIAM J. on Numerical Analysis*, vol. 41, no. 1, pp. 325–363, 2003.

[24] J.-M. Mirebeau, "Anisotropic Fast-Marching on Cartesian Grids Using Lattice Basis Reduction," *SIAM J. on Numerical Analysis*, vol. 52, no. 4, pp. 1573–1599, 2014.

[25] ——, "Fast Marching methods for Curvature Penalized Shortest Paths," *J. of Math. Imaging and Vision*, vol. 60, no. 6, pp. 784–815, 2018.

[26] M. Bardi and I. Capuzzo-Dolcetta, *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*. Birkhäuser Boston, MA, 00 1997.

[27] J. Mirebeau *et al.*, "Massively parallel computation of globally optimal shortest paths with curvature penalization," *Concurrency and Computation: Practice and Experience*, vol. 35, no. 2, p. e7472, 2023.

[28] F. Zhu and P. J. Antsaklis, "Optimal control of hybrid switched systems: A brief survey," *Discrete Event Dynamic Systems*, vol. 25, no. 3, pp. 345–364, 2015.

[29] A. S. Bortakovskii, "Sufficient optimality conditions for controlling switched systems," *Journal of Computer and Systems Sciences International*, vol. 56, no. 4, pp. 636–651, 2017.

[30] E. Heiden *et al.*, "Bench-MR: A motion planning benchmark for wheeled mobile robots," *IEEE Robot. and Autom. Lett.*, vol. 6, no. 3, pp. 4536–4543, 2021.

[31] S. M. LaValle, "Rapidly-Exploring Random Trees: A New Tool for Path Planning," *The annual research report*, 1998.

[32] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The Int. J. of Robot. Research*, vol. 30, no. 7, pp. 846–894, 2011.

[33] J. D. Gammell *et al.*, "Informed RRT*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2014, pp. 2997–3004.

[34] Y. Li *et al.*, "Asymptotically optimal sampling-based kinodynamic planning," *The Int. J. of Robot. Research*, vol. 35, no. 5, pp. 528–564, 2016.

[35] ROS Navigation Stack. [Online]. Available: https://github.com/ros-planning/navigation