Calibrated Ensemble-Based Chance-Constrained Safety for Autonomous Mobile Robots in Lean Intralogistics under Localisation Uncertainty

Toheeb Babatunde Ibrahim[1*], Olasubomi Akanbi[2], Joseph Oladejo[1], Warith Akinade[1]

[1]Department of Mathematics and Statistics, Bowling Green State University, Bowling Green, OH 43403, USA
[2]Department of Systems Engineering, Bowling Green State University, Bowling Green, OH 43403, USA

Abstract

Autonomous mobile robots that run lean milk-run routes must keep a safe separation from operators, pallets and aisle boundaries while meeting a fixed delivery cycle. Their protective fields are usually sized once, for the worst localisation in the plant, so the vehicle stops at obstacles it could pass wherever it is well localised. We study uncertainty-aware safety filters for this setting, in particular an ensemble-based chance constraint that evaluates a discrete-time control barrier function on the members of an ensemble transform Kalman filter. Under a stated Gaussian assumption, the filter bounds the one-step probability that each separation requirement fails. We show that the constraint acts as a protective field whose margin is proportional to the localisation uncertainty in the direction of each barrier, and derive from it a local necessary condition on the pallet protrusion the vehicle can pass. Protective stops are applied only when a feasibility test shows that no certified command exists. In 4680 simulated runs of a four-station milk run with intruding pallets and crossing operators, a fixed margin sized for an occluded section gave a mean cycle time of 169.2 s, 66% of cycles on time and 17.2 operator calls per hour. The ensemble filter gave 127.8 s, 81% and 5.8 calls per hour, but 3.6 more separation events per hour, with 9 contacts in 10.7 simulated hours against 3 in 14.1 hours. At matched safety, all uncertainty-aware filters shared one frontier well below the fixed margin, and the ensemble filter showed no advantage over a linearised Kalman-filter chance constraint. The chance constraints differ mainly in calibration: their realised one-step risk followed the design level, whereas that of a covariance-scaled margin did not. The results are from simulation; the code, the data and a SymPy verification of the derivations accompany the paper.

Keywords: Lean manufacturing, intralogistics, autonomous mobile robots, control barrier functions, ensemble transform Kalman filter, chance constraints, localisation uncertainty

Journal of Science, Technology and Social Transformation (2026) Doi: 10.64235/A934jf02

1. Introduction

1.1. Background

Lean production supplies assembly lines in small lots at short, fixed intervals, so that line-side stock stays low and problems surface quickly [1–3]. In-plant milk runs implement this principle: a vehicle follows a fixed route on a fixed cycle, delivers bins to line-side stations and returns empties [4–8]. The route cycle is tied to the takt of the line, and a late vehicle either starves a station or forces the plant to hold more line-side stock, which is the waste the milk run was introduced to remove. Autonomous mobile robots (AMRs) and automated guided vehicles increasingly

Corresponding Author: Toheeb Babatunde Ibrahim, Department of Mathematics and Statistics, Bowling Green State University, Bowling Green, OH 43403, USA, Email: toheeb.ibrahim@aims-cameroon.org

How to cite this article: Ibrahim, T. B., Akanbi, O., Oladejo, J., & Akinade, W. (2026). Calibrated Ensemble-Based Chance-Constrained Safety for Autonomous Mobile Robots in Lean Intralogistics under Localisation Uncertainty. *Journal of Science, Technology and Social Transformation* 2(4), 26–41.

Source of support: Nil

Conflict of interest: None

drive these routes [9–11], and digital technologies of this kind are widely regarded as enablers of lean practice [12–14].

An AMR on a milk run shares its aisle with operators, pallets and other vehicles, so its safety system determines how much of the planned cycle is achieved in operation. Industrial vehicles protect themselves with protective fields and separation distances fixed at commissioning [15, 16]. A field that is too small risks contact; a field that is too large makes the vehicle slow down or stop for obstacles it could have passed, and a vehicle that stops in an aisle blocks it until an operator intervenes. In lean terms, unnecessary stops are waiting, and operator call-outs are a recurring abnormality.

1.2. Problem and Research Gap

The separation a robot actually needs depends on how well it knows where it is. Localisation quality varies along a route: reflectors and features are dense in some areas and occluded in others, for instance by stored material, and between fixes the estimate drifts with the odometry [10, 17]. A margin sized for the worst section of the plant is too large everywhere else, and a margin sized for the typical section is too small where localisation is poor.

Control barrier functions (CBFs) provide safety filters that modify a nominal command as little as possible while keeping the robot in a safe set [18–21]. In manufacturing they have been used to keep collaborative manipulators within ISO separation and speed limits while preserving throughput [22, 23], for mobile manipulators and AMRs working near people [24, 25], and for an industrial mobile robot under localisation uncertainty with a worst-case error bound [26]. Robust formulations of this kind tighten the barrier by a bound on the estimation error, either fixed [26–28] or adapted online from the estimator covariance [29]. Probabilistic formulations instead bound the probability that the barrier condition fails, using the belief of a Kalman filter or of a particle set [30–34]. The ensemble-transform CBF (ET-CBF) of our companion work [35] evaluates the barrier condition on the members of an ensemble transform Kalman filter (ETKF) [36, 37] and enforces a chance constraint without Jacobians.

Probabilistic barrier constraints under uncertain sensing [32, 38] and adaptive robust CBFs with several obstacles and hardware experiments [29] are therefore established. Three questions that matter for lean intralogistics remain open. First, these filters have not been studied on a production route, where several barriers are active at once, operators cross the aisle, the vehicle stops at stations, and a rule is needed for what the vehicle does when no command can be certified and for how such a stop is resolved. Second, safety filters are assessed by violation rates and tracking error, whereas a plant measures cycle time, cycle adherence, deliveries per hour and operator call-outs; how the risk level and the localisation quality translate into these quantities has not been reported. Third, to our knowledge no study compares fixed-margin, adaptive-margin and calibrated chance-constrained filters on the same intralogistics task at matched safety.

1.3. Contributions

The contribution of this paper is applied: it integrates uncertainty-aware safety filtering with the evaluation of a lean milk run. The filter builds on known constructions [31, 32, 35]; we restate the parts we use in a self-contained form, with proofs, and state the assumptions under which each guarantee holds. The contributions are as follows.

1. *Intralogistics formulation.* We formulate the safety filter for a one-way milk-run loop with aisle boundaries, line-side pallets that intrude into the aisle, operators crossing at crosswalks, and station stops, and state the per-barrier and joint one-step risk of the filter under explicit Gaussian assumptions (Proposition 1).
2. *Protective-field interpretation.* We show that the chance constraint is a deterministic barrier condition with a margin proportional to the localisation uncertainty in the direction of each barrier (Proposition 2). This gives a closed-form, local necessary condition for passing a pallet: the largest protrusion at which the filter can certify a command decreases with the lateral localisation standard deviation, asymptotically at a rate of 2κ (Corollary 1). The condition links localisation quality to the tolerance of a route for 5S deviations in the aisle.
3. *Stop-and-go operation.* We specify a protective-stop rule that is triggered only when a feasibility test shows that no certified command exists, together with an andon rule, and we show why multiplicative covariance inflation, used in the companion work, makes the ensemble diverge when the vehicle stands at a station; relaxation to prior spread avoids this divergence (Section 4.5).
4. *Evaluation in lean terms.* In 4680 simulated runs of a milk-run loop, covering 435 hours of operation, we compare ET-CBF with a certainty-equivalent filter, a fixed margin sized at commissioning, a covariance-adaptive margin, a linearised EKF

chance constraint and a true-state reference. We report flow and safety side by side, trace the frontier between them for each family of filters, and test the sensitivity of the results to the operator response and the recovery procedure. At matched safety, the uncertainty-aware filters share one frontier, and the ensemble filter shows no advantage over its linearised counterpart in this setting. We verify the derivations and their implementation symbolically with SymPy and by sampling (Appendix A), and release the code and data.

1.4. Organisation

Section 2 reviews related work. Section 3 states the problem. Section 4 develops the method. Section 5 describes the experimental design, and Section 6 reports the results. Section 7 discusses implications and limitations, and Section 8 concludes.

2. Related Work

2.1. Lean Part Supply and Mobile Robots

Part feeding for mixed-model assembly is organised around supermarkets and tow-train or milk-run routes with fixed cycles [4–7]; operating strategies for in-plant milk runs are compared in [8]. Reviews of AMRs in intralogistics identify decentralised navigation, human interaction and robustness as open planning and control issues [9, 10], and warehouse-oriented surveys treat fleet coordination at length [39–41]. The literature on lean and Industry 4.0 argues that automation supports lean flow when it is reliable and transparent [12–14]; built-in quality and the stopping of a process at the first abnormality (jidoka, andon) are central lean principles [2, 3, 42]. These works do not model how the safety layer of a vehicle affects cycle adherence.

2.2. Safety of Industrial Mobile Robots

Safety requirements for driverless industrial trucks are set by ISO 3691-4 [15], and the risks of AGV systems are reviewed in [16]. Human-aware navigation stacks for industrial vehicles combine people tracking with planning and reactive safety layers [43], and infrastructure-based AMRs have been deployed in a truck plant, where users named safety among their main concerns [44]. In manufacturing cells, CBFs have been used to satisfy ISO/TS 15066 and ISO 10218 limits while retaining robot speed [22, 23, 45], and for mobile robots among people [24, 25]. These filters assume that the relevant states are measured accurately.

2.3. Safety Filters under Estimation Uncertainty

Robust CBFs tighten the barrier by an error bound. Measurement-robust CBFs do so for perception error [27, 28]; Miksits et al. applied a robust CBF to an industrial mobile robot whose localisation error is bounded by a constant [26]; and Daş et al. adapt the parameters of a robust CBF online to reduce its conservatism, merge several obstacles into one numerical barrier, and demonstrate the approach on a tracked vehicle [29]. Chance-constrained planning under Gaussian uncertainty goes back to [46]. Stochastic and belief-space CBFs bound the probability of violating the barrier condition with an extended Kalman filter (EKF) [30–32, 34, 47], with particle beliefs [33], or with sigma points in policy optimisation [48]; discrete-time analyses quantify the accumulated risk [49, 50], and Luo et al. derive probabilistic safety barrier certificates for multi-robot collision avoidance under uncertain sensing [38]. CBF filters can also stop a robot at an undesired equilibrium, a behaviour analysed as deadlock in [51, 52]; in an aisle this appears as a blocked vehicle.

2.4. Ensemble Filtering

Ensemble Kalman filters propagate a sample of states through the nonlinear model [53]; the ETKF updates the sample with a deterministic square-root transform [36, 37, 54, 55]. Finite ensembles underestimate the error covariance, which is commonly corrected by multiplicative inflation or by relaxing the analysis spread towards the prior spread [56]. The companion manuscript [35] couples an ETKF with a CBF for a single obstacle. Because that manuscript is not yet published, Section 4 restates the parts of the construction used here, with proofs, and none of the results below depends on it. Table 1 positions the present work.

Table 1: Positioning of this work among safety filters for mobile robots under localisation uncertainty.

Approach	Margin	Guarantee	Setting evaluated	Performance measures
Fixed-bound robust CBF [26–28]	Constant, worst case	Deterministic under the error bound	Single robot, few obstacles	Clearance, tracking

Approach	Margin	Guarantee	Setting evaluated	Performance measures
Online adaptive robust CBF [29]	Adapted online to reduce conservatism	Deterministic under the error bound	Tracked vehicle, multiple obstacles, hardware	Conservatism, feasibility, control effort
Probabilistic safety barrier certificates [38]	Chance constraint under Gaussian sensing noise	Probabilistic, per step	Multiple robots	Collision avoidance
EKF or particle belief CBF [31–34]	Directional, linearised or sampled	One-step risk under the belief	Single robot, short manoeuvres	Violation rate
ET-CBF [35]	Directional, ensemble	One-step risk under Gaussian assumptions	Single obstacle	Violation rate, calibration
This work	**Directional, ensemble or linearised; several and moving barriers**	**Per-barrier one-step risk under stated assumptions; union bound; feasibility-tested protective stop**	**Simulated milk-run loop with pallets, operators and stations**	**Cycle time, adherence, deliveries, call-outs, infringements, contacts**

3. Problem Setting

3.1. Milk-Run Loop and Robot Model

We consider one AMR that supplies four line-side stations on a one-way loop around an assembly line (Fig. 1). The centreline of the loop consists of two straights of length L joined by two semicircles of radius R; equivalently, it is the set of points at distance R from the segment that joins the two semicircle centres, which we call the spine. We denote by $d_{\mathrm{sp}}(q)$ the distance of a point $q \in \mathbb{R}^2$ to the spine. The robot stops at each station for a dwell time T_d to exchange bins and then continues. The planned cycle time of the route is T_c.

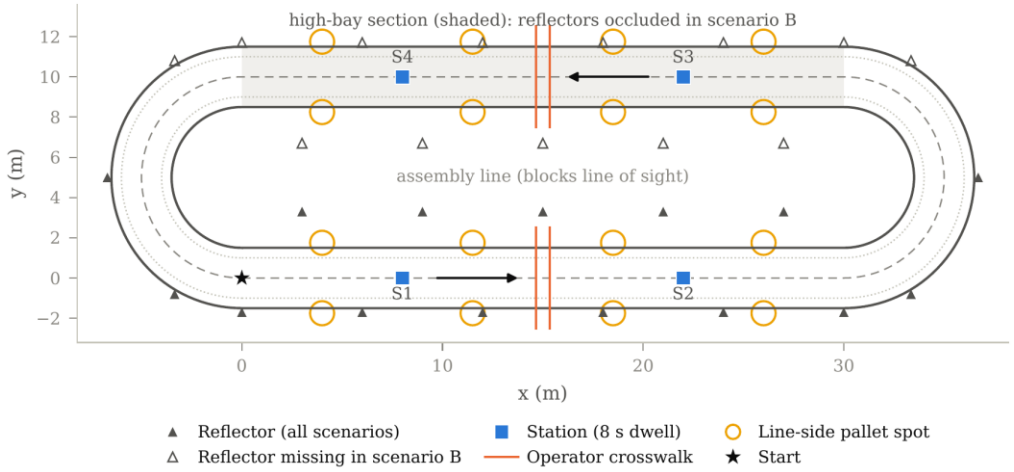


Figure 1: Milk-run loop used in the study. The robot drives anticlockwise from the start, stops at stations S1–S4, and shares the aisle with operators at two crosswalks and with pallets that may intrude from eight line-side spots. Dashed line: centreline; dotted lines: corridor for the reference point of the robot.

The robot is a differential-drive vehicle. As in [35], we describe it by a reference point q located a distance ℓ ahead of the wheel axle and by its heading θ, so that the state is $x = (q, \theta) \in \mathbb{R}^3$ and the input $u = (v, \omega)$ collects the forward speed v (m/s) and the turn rate ω (rad/s). With control period Δt, the sampled model is

$$x_{k+1} = x_k + G(x_k)u_k + w_k, \qquad G(x) = \Delta t \begin{bmatrix} \cos\theta & -\ell\sin\theta \\ \sin\theta & \ell\cos\theta \\ 0 & 1 \end{bmatrix}, \tag{1}$$

where k is the time step and $w_k \sim \mathcal{N}(0, Q)$ is process noise that represents odometry error, with $Q = \mathrm{diag}(\sigma_w^2, \sigma_w^2, \sigma_\theta^2)$ while the robot moves and $Q = 0$ at standstill. The input is bounded, $u_k \in \mathcal{U} = [0, v_{\max}] \times [-\omega_{\max}, \omega_{\max}]$; the vehicle cannot reverse. We write $G_q(\theta) \in \mathbb{R}^{2\times 2}$ for the first two rows of G, which map the input to the displacement of q.

3.2. Localisation

The robot measures the range and bearing of wall-mounted reflectors with a scanner at the reference point,

$$y_k^{(i)} = \begin{bmatrix} \| l_i - q_k \| \\ \mathrm{atan2}(l_i - q_k) - \theta_k \end{bmatrix} + v_k^{(i)}, \tag{2}$$

where $l_i \in \mathbb{R}^2$ is the position of reflector i and $v_k^{(i)} \sim \mathcal{N}\big(0, \mathrm{diag}(\sigma_r^2, \sigma_b^2)\big)$ is measurement noise with range and bearing standard deviations σ_r (m) and σ_b (rad). A reflector

is measured only if it lies within the sensor range r_s and the line of sight does not cross the assembly line inside the loop. Where no reflector is visible, the estimate is propagated by the odometry model (1) alone and its uncertainty grows. An estimator provides, at every step, a belief about x_k given the measurement history Y_k: an EKF provides a mean $\hat{x}_k$ and covariance P_k, and an ETKF provides N members $x_k^{(1)}, \dots, x_k^{(N)}$.

3.3. Safety Requirements

The robot body is covered by a disc of radius ρ_r centred at q. Safety requirements are expressed for q through barrier functions h_j that are nonnegative when requirement j holds. All barriers used here are signed distances in metres, so that their gradient with respect to q, which we denote by $n_j(q)$, is a unit vector.

- *Aisle boundaries.* The reference point must stay within a lateral distance w of the centreline, where w is the half-width of the aisle minus ρ_r. The outer and inner boundaries give

$$h_{\text{out}}(q) = (R + w) - d_{\text{sp}}(q), \qquad h_{\text{in}}(q) = d_{\text{sp}}(q) - (R - w). \tag{3}$$

- *Pallets.* A pallet with centre o_j and radius ρ_p must be kept at the required separation d_{sep}, so that

$$h_j(q) = \| q - o_j \| - r_j, \qquad r_j = \rho_p + \rho_r + d_{\text{sep}}. \tag{4}$$

- *Operators.* An operator who crosses the aisle is a disc of radius ρ_h whose centre $o_{j,k}$ moves by a known displacement $\Delta o_{j,k} = o_{j,k+1} - o_{j,k}$ per step; the barrier is (4) with $r_j = \rho_h + \rho_r + d_{\text{sep}}$. We assume that operators are tracked by plant infrastructure, as in [44], so that the only uncertainty in the barrier is the position of the robot itself.

The set of barriers that are relevant at step k, namely the two aisle boundaries and the pallets and operators near the estimated position, is denoted by $\mathcal{J}_k$. Following the discrete-time CBF condition [49, 57], each barrier must satisfy

$$h_j(q_{k+1}, o_{j,k+1}) \geq (1 - \gamma)\, h_j(q_k, o_{j,k}), \tag{5}$$

where $\gamma \in (0,1]$ is the rate at which the robot may approach a boundary. If (5) holds at every step and $h_j \geq 0$ initially, requirement j is never violated.

3.4. Lean Performance Measures

We evaluate a safety filter by the quantities a plant would track.

- *Cycle time*: the time for one loop including station dwells; we report its mean and 95th percentile. The *free-flow cycle time* T_{free} is the cycle time without pallets and operators.
- *Cycle adherence*: the share of cycles completed within the planned cycle time T_c.
- *Deliveries per hour*: station visits per hour of operation.
- *Protective-stop share*: the share of travel time during which the filter cannot certify any command and the vehicle stands still.
- *Andon calls per hour*: the rate at which a blocked vehicle requires an operator.
- *Separation infringements per hour*: events in which the true separation to a pallet or an operator falls below d_{sep}; an infringement that reaches zero separation is a *contact*. *Lane departures* are events in which the reference point leaves the corridor.
- *Realised one-step risk*: among the steps and barriers at which the filter constraint is tight, the share at which the true state violates (5). For a calibrated filter this equals the design risk level of each barrier. The *joint one-step failure rate* is the share of certified steps at which any barrier in $\mathcal{J}_k$ is violated.

3.5. Problem Statement

Given the belief of the estimator, a nominal command u_{nom} from the route controller and a risk level $\varepsilon \in (0,1)$, we seek at every step the command closest to u_{nom} such that each barrier condition (5) fails with probability at most ε under the belief, together with a rule for the steps at which no such command exists. We then ask how ε and the localisation quality determine the lean performance measures, and how the answer compares with margins fixed at commissioning.

4. METHOD

4.1. Ensemble-Transform CBF for Distance Barriers

We first state the ensemble construction, following [35], for the barriers of Section 3. The condition (5) is nonlinear in the input. For barrier j we therefore define

$$\psi_j(x, u, w) = \gamma h_j(q) + n_j(q)^{\top}\big(G_q(\theta)u + w_q - \Delta o_j\big), \tag{6}$$

where $w_q \in \mathbb{R}^2$ is the position part of the process noise and $\Delta o_j = 0$ for static barriers.

Lemma 1 (Tangent condition). *Let h_j be the pallet or operator barrier (4) or the inner boundary h_{in}. Then*

$\psi_j(x_k, u_k, w_k) \geq 0$ *implies the barrier condition (5). For the outer boundary* h_{out}, $\psi_j \geq 0$ *implies (5) up to a residual of at most* $\| q_{k+1} - q_k \|^2/\big(2(R-w)\big)$.

Proof. The barriers (4) and h_{in} are convex functions of the relative position $q - o_j$, because the distance to a point and the distance to a segment are convex. Convexity gives $h_j\big(q_{k+1}, o_{j,k+1}\big) \geq h_j\big(q_k, o_{j,k}\big) + n_j^{\top}\big(q_{k+1} - q_k - \Delta o_j\big)$, and substituting the model (1) yields $h_j\big(q_{k+1}, o_{j,k+1}\big) - (1-\gamma)h_j\big(q_k, o_{j,k}\big) \geq \psi_j$. The outer boundary is the negative of a convex distance whose Hessian has norm at most $1/d_{\text{sp}}(q) \leq 1/(R-w)$ inside the corridor, so the first-order expansion holds up to the stated second-order residual. □

With the speeds and dimensions used here, the residual of the outer boundary is at most 3.3 mm per step at the input bounds and 1.8 mm at cruise speed (Appendix A); we count it against the filter in the evaluation.

Since ψ_j is affine in u and in w for a fixed state, its mean and variance over the belief are available from the ensemble. For every member $x^{(i)} = \big(q^{(i)}, \theta^{(i)}\big)$ we compute

$$a_{ij} = \gamma h_j\big(q^{(i)}\big) - n_j\big(q^{(i)}\big)^{\top}\Delta o_j, \qquad b_{ij} = G_q\big(\theta^{(i)}\big)^{\top} n_j\big(q^{(i)}\big), \tag{7}$$

so that $\psi_j\big(x^{(i)}, u, w\big) = a_{ij} + b_{ij}^{\top}u + n_j^{\top}w_q$. We denote by $\bar{a}_j$ and $\bar{b}_j$ the ensemble means of a_{ij} and b_{ij}, and by $a_j' \in \mathbb{R}^N$ and $B_j' \in \mathbb{R}^{N\times 2}$ the corresponding deviations from the mean. Because n_j is a unit vector, the process noise contributes the variance σ_w^2. The ensemble estimates of the mean and variance of ψ_j are therefore

$$\hat{\mu}_j(u) = \bar{a}_j + \bar{b}_j^{\top}u, \qquad \hat{\sigma}_j^2(u) = \frac{\| a_j' + B_j'u \|^2}{N-1} + \sigma_w^2. \tag{8}$$

The filter solves the second-order cone program

$$\begin{aligned} \min_{u\in\mathcal{U},\, s\geq 0} \quad & \| S(u - u_{\text{nom}}) \|^2 + \lambda\, \mathbf{1}^{\top}s \\ \text{s.t.} \quad & \hat{\mu}_j(u) + s_j \;\geq\; \kappa\, \hat{\sigma}_j(u), \qquad j \in \mathcal{J}_k, \end{aligned} \tag{9}$$

where $S = \text{diag}(1, \ell)$, so that the cost is the squared deviation of the velocity of the reference point, $s_j \geq 0$ is a slack variable with penalty λ, and κ is the constraint gain. The Gaussian gain for a risk level ε is $\kappa = \Phi^{-1}(1-\varepsilon)$, with Φ the standard normal distribution function. With a small ensemble, $\hat{\sigma}_j$ is itself a noisy estimate, and we use the inflated gain

$$\kappa_N(\varepsilon) = t_{N-1,\,1-\varepsilon}\sqrt{1 + 1/N}, \tag{10}$$

where $t_{N-1,1-\varepsilon}$ is the $(1-\varepsilon)$ quantile of Student's t distribution with $N-1$ degrees of freedom. If $\psi^{(1)}, \ldots, \psi^{(N)}$ and ψ were independent draws from one Gaussian distribution, $(\psi - \bar{\psi})/\big(s\sqrt{1+1/N}\big)$, with $\bar{\psi}$ and s the sample mean and standard deviation, would follow that t distribution, and (10) would give exactly the risk ε. ETKF members are neither independent nor draws from the posterior, so we use (10) as a heuristic correction for sampling error, not as a coverage guarantee; Experiment 4 tests it. A thin QR factorisation of $\big[a_j'\ B_j'\big]$ reduces each cone to dimension five, so the size of (9) depends on the number of active barriers and not on N.

4.2. Several and Moving Barriers

In an aisle, several barriers are active at once and some of them move. The following assumption makes explicit what the risk statement requires of the estimator.

Assumption 1 (Ensemble consistency). Conditioned on the measurement history Y_k, the mean and variance of $\psi_j(x_k, u_k, w_k)$ equal $\hat{\mu}_j(u_k)$ and $\hat{\sigma}_j^2(u_k)$ for every $j \in \mathcal{J}_k$, and ψ_j is conditionally Gaussian.

Proposition 1 (Per-barrier and joint one-step risk). *Let Assumption 1 hold, and let the command* u_k *satisfy the constraints of problem (9) with* $s_j = 0$ *and gain* $\kappa = \Phi^{-1}\big(1-\varepsilon_j\big)$ *for each* $j \in \mathcal{J}_k$. *Then each pallet, operator and inner-boundary barrier fails (5) at step* k *with conditional probability at most* ε_j, *and*

$$\mathbb{P}[\exists\, j \in \mathcal{J}_k\colon\ h_j\big(q_{k+1}, o_{j,k+1}\big) < (1-\gamma)h_j\big(q_k, o_{j,k}\big)] \leq \sum_{j\in\mathcal{J}_k} \varepsilon_j\,, \tag{11}$$

and the same bound holds for the outer boundary up to the residual of Lemma 1.

Proof. Under Assumption 1, ψ_j is conditionally distributed as $\mathcal{N}\left(\hat{\mu}_j(u_k), \hat{\sigma}_j^2(u_k)\right)$. With $s_j = 0$, the constraint gives $\hat{\mu}_j \geq \kappa\hat{\sigma}_j$, so $\mathbb{P}\big[\psi_j < 0 \mid Y_k\big] = \Phi\big(-\hat{\mu}_j/\hat{\sigma}_j\big) \leq \Phi(-\kappa) = \varepsilon_j$. By Lemma 1, the failure of barrier j implies $\psi_j < 0$, which gives the per-barrier bound, and the union bound over j gives (11). □

Proposition 1 is conditional, and three limits apply to its use. First, it holds only to the extent that Assumption 1 holds; a finite ETKF satisfies it only approximately, and Section 7 reports how far it holds in our runs. Second, with the same level ε for every barrier, the joint bound is $|\mathcal{J}_k|\varepsilon$ and not ε; a plant that needs a joint level ε has to set $\varepsilon_j = \varepsilon/|\mathcal{J}_k|$. We use the same level for every barrier, so with 2.6 barriers in $\mathcal{J}_k$ on average in our runs, the joint one-step bound at $\varepsilon = 0.05$ is about 13%, and we report the realised joint failure rate alongside the per-barrier risk. Third, the bound concerns one step. Localisation errors persist over many steps, so the probability of a separation event over a

cycle or a route is not bounded by ε; we measure it directly as the rate of infringements and contacts per hour.

4.3. The Chance Constraint as a Protective Field

Industrial vehicles are commissioned with protective fields, and the fixed-margin and adaptive-margin filters of Section 4.6 tighten every barrier by a distance. The chance constraint has the same form, with a margin that follows the uncertainty in the direction of the barrier.

Proposition 2 (Uncertainty-proportional margin). *For a static barrier j, the constraint of problem (9) with $s_j = 0$ is identical to*

$$\gamma\left(\bar{h}_j - m_j(u)\right) + \bar{b}_j^{\top} u \;\geq\; 0, \qquad m_j(u) = \frac{\kappa}{\gamma}\,\hat{\sigma}_j(u), \tag{12}$$

where $\bar{h}_j$ is the ensemble mean of $h_j\left(q^{(i)}\right)$. The margin satisfies

$$\frac{\kappa\,\sigma_w}{\gamma} \;\leq\; m_j(u) \;\leq\; m_j^0 + \frac{\kappa}{\gamma}\frac{\|\,B_j' u\,\|}{\sqrt{N-1}}, \qquad m_j^0 = \kappa\sqrt{s_{h,j}^2 + \sigma_w^2/\gamma^2}, \tag{13}$$

where $s_{h,j}^2$ is the sample variance of $h_j\left(q^{(i)}\right)$ over the ensemble. To first order in the ensemble spread, $s_{h,j}^2 = n_j^{\top}\hat{P}_q\, n_j$, where $\hat{P}_q$ is the sample covariance of the position: the variance of the position error in the direction of the barrier normal.

Proof. For a static barrier, $\bar{a}_j = \gamma\bar{h}_j$, and (12) is a rearrangement of $\hat{\mu}_j(u) \geq \kappa\hat{\sigma}_j(u)$. The lower bound in (13) follows from $\hat{\sigma}_j \geq \sigma_w$. For the upper bound, $\hat{\sigma}_j(u)$ is the Euclidean norm of the vector $\left(\left(a_j' + B_j' u\right)/\sqrt{N-1},\ \sigma_w\right)$, and the triangle inequality applied to its two parts $\left(a_j'/\sqrt{N-1}, \sigma_w\right)$ and $\left(B_j' u/\sqrt{N-1}, 0\right)$, together with $\|\,a_j'\,\|^2/(N-1) = \gamma^2 s_{h,j}^2$, gives the result. The last statement follows from $h_j\left(q^{(i)}\right) = h_j(\bar{q}) + n_j^{\top}\left(q^{(i)} - \bar{q}\right)$ to first order. □

The term m_j^0 is the stand-off distance: at rest, the robot is certified with respect to barrier j if and only if the mean barrier value is at least m_j^0. The second term of the upper bound in (13) is the contribution of the heading spread, which grows with speed. In comparison, a fixed margin $\bar{\epsilon}$ is the same everywhere on the route, and an isotropic adaptive margin $c\sqrt{\lambda_{\max}\left(P_q\right)}$, where $\lambda_{\max}$ is the largest eigenvalue, is at least $c\sqrt{n_j^{\top} P_q n_j}$ for every barrier, so it charges the uncertainty of the worst direction to all barriers.

Corollary 1 (Necessary passing width). *Consider a straight aisle whose corridor for the reference point has width $2w$, obstructed by a static obstacle that protrudes a distance δ into the corridor, so that the free width is $g = 2w - \delta$. Suppose that the heading spread of the ensemble is negligible and that the lateral position has sample standard deviation s_e. To first order in the spread, the filter can certify a command at the narrowest cross-section only if*

$$\delta \;\leq\; \delta_{\max} = 2w - 2\kappa\sqrt{s_e^2 + \sigma_w^2/\gamma^2}. \tag{14}$$

For a filter with a fixed margin $\bar{\epsilon}$, the corresponding threshold is $\delta_{\max} = 2w - 2\bar{\epsilon}$.

The condition is local and necessary. It does not guarantee that the robot completes the manoeuvre, which also depends on the approach, the heading and the evolution of the estimate during the pass.

Proof. At the narrowest cross-section, the normals of the obstacle barrier and of the opposite aisle boundary are opposed, so with a common heading $\bar{b}_1 = -\bar{b}_2$ and $B_1' = B_2' = 0$. Adding the two constraints (12) gives $\bar{h}_1 + \bar{h}_2 \geq m_1^0 + m_2^0 = 2\kappa\sqrt{s_e^2 + \sigma_w^2/\gamma^2}$, and $\bar{h}_1 + \bar{h}_2 = g$ to first order when the ensemble mean lies on that cross-section. □

Equation (14) quantifies a trade-off that lean plants manage implicitly. The threshold $\delta_{\max}$ is an upper limit on the tolerance of the route to 5S deviations such as a pallet set down outside its floor marking. Each centimetre of lateral localisation standard deviation lowers that limit by up to 2κ centimetres, about 3.3 cm at $\varepsilon = 0.05$, and a margin fixed for the worst section of the plant lowers it on the whole route.

4.4. Protective Stop and Andon

A positive slack in the solution of problem (9) does not by itself show that no certified command exists, because with a finite penalty λ the program may trade a small slack against tracking cost. Whenever a slack exceeds the solver tolerance of 10^{-6}, we therefore solve problem (9) a second time with λ replaced by $\lambda_f = 10^9$. On $\mathcal{U}$ the tracking cost is at most $v_{\max}^2 + 4\ell^2\omega_{\max}^2 < 4$, so any zero-slack command has a lower cost than a slack of 4×10^{-9}, and the second program returns a zero-slack command whenever one exists, up to the solver tolerance. If it does, that command is applied. Otherwise no command can be certified, and moving on a slack-optimal command would add odometry error to a belief that is already too uncertain. The vehicle then makes a *protective stop*, $u_k = 0$. In Experiment 2, the second program was solved at 204,026 steps and found a zero-slack command at 13 of them. Because the process noise in (1) vanishes at standstill, a protective stop cannot infringe a static barrier; only an operator who walks into

the stationary vehicle can reduce the separation. If the vehicle makes no progress for a time T_a while it is neither at a station nor yielding to an operator, it raises an *andon call*. In the base case, an operator arrives after a response time T_r, removes the pallets within reach of the vehicle, places the vehicle back on the centreline and re-initialises its localisation. Experiment 6 varies T_r and replaces this recovery by one in which the operator first only re-localises the vehicle and intervenes physically only if it calls again at the same spot. Protective stops and andon calls are logged with their position on the route, which makes poorly localised sections visible (Section 6.3). The same rule is applied to all filters in the comparison.

4.5. Ensemble Filtering in Stop-and-Go Operation

The companion work inflates the ETKF covariance multiplicatively by a factor $\rho > 1$ at every update. This is adequate for a robot that keeps moving, but a milk-run vehicle stands at stations. At standstill the forecast adds no spread, while every update still multiplies the variance of directions that the sensors do not observe by ρ. At an update rate f, the variance in such a direction therefore grows by ρ^{fT_d} during a dwell: a factor of about 50 for $\rho = 1.05$, $f = 10\,\text{Hz}$ and $T_d = 8\,\text{s}$. In a pilot run with sparse reflectors, the lateral standard deviation of the ensemble exceeded 0.5 m after a few stops and the vehicle did not complete a cycle.

We therefore replace multiplicative inflation by relaxation to prior spread [56]. After each ETKF update, the deviations of every state variable are rescaled by

$$1 + \alpha \, \frac{s^b - s^a}{s^a}, \tag{15}$$

where s^b and s^a are the ensemble standard deviations of that variable before and after the update and $\alpha \in [0,1]$ is the relaxation coefficient. The analysis spread is never larger than the prior spread, so the correction cannot diverge at standstill. We selected $\alpha = 0.3$ in a pilot study without obstacles: the median normalised position error squared was between 1.84 and 2.04 in the three scenarios of Section 5 (the consistent value is 2), against 2.45 to 2.62 without relaxation, and the share of lateral errors beyond 1.645 estimated standard deviations on one side was 4.8% to 5.8% (nominal 5%), against 6.8% to 8.6%.

4.6. Comparators

All filters minimise the cost of problem (9), use the same barriers and input bounds, and apply the protective stop of Section 4.4. They differ in the constraint.

- *Certainty equivalent (CE).* The deterministic condition $\gamma h_j(\hat{q}) + n_j^{\top}\big(G_q(\hat{\theta})u - \Delta o_j\big) \geq 0$ on the EKF mean.
- *Fixed margin.* The CE condition with h_j replaced by $h_j - \bar{\epsilon}$, where $\bar{\epsilon}$ is a bound on the position error fixed at commissioning, in the spirit of [26].
- *Adaptive margin.* The same with $\bar{\epsilon}_k = c\sqrt{\lambda_{\max}\big(P_{q,k}\big)}$, where $P_{q,k}$ is the position block of the EKF covariance; $c = 2.45$ corresponds to a 95% confidence region in two dimensions. This covariance-scaled margin is a simple stand-in for adaptive robust CBFs; it does not reproduce the online parameter adaptation of [29].
- *CC-EKF.* The Gaussian chance constraint with mean and variance of ψ_j obtained by linearising ψ_j at the EKF mean, representing [31, 32, 34].
- *ET-CBF.* Problem (9) with the ETKF of Section 4.5 and the gain (10).
- *Oracle.* The chance constraint on the true state, with the process noise as the only uncertainty; it is not implementable and serves as a reference.

5. Experimental Design

5.1. Plant and Robot

Table 2 lists the parameters. The aisle is 3 m wide, the loop is 91.4 m long, and the robot cruises at 1.2 m/s. The route controller follows the centreline with a look-ahead of 1 m on the estimated position, decelerates towards the next station and dwells there; its command is the nominal input u_{nom}. The free-flow cycle time, measured over 30 cycles without pallets and operators, is $T_{\text{free}} = 113.3\,\text{s}$ (standard deviation 0.4 s), and the planned cycle time is $T_c = 120\,\text{s}$, which leaves an allowance of about 6%.

In every run, each of the eight line-side spots holds a pallet with probability 0.5. A pallet protrudes into the corridor of the reference point by a distance drawn uniformly from 0.3 to 1.5 m, from the inner or the outer side with equal probability; protrusions above 1 m reach beyond the centreline, so the robot must deviate. Operators arrive at each crosswalk with exponentially distributed gaps of mean 20 s and cross at 1.3 m/s. An operator does not step into the aisle while the robot is between 3 m before and 1.5 m after the crosswalk, but otherwise does not react to it.

Table 2: Parameters of the simulation study.

Quantity	Value
Straight length L; end radius R	30 m; 5 m
Corridor half-width w; robot radius ρ_r	1.0 m; 0.5 m

Quantity	Value
Pallet radius ρ_p; operator radius ρ_h	0.6 m; 0.3 m
Required separation d_{sep}	0.1 m
Control period Δt; look-ahead ℓ	0.1 s; 0.3 m
Cruise speed; $v_{\max}$; $\omega_{\max}$	1.2 m/s; 1.5 m/s; 2 rad/s
Process noise σ_w; σ_θ (per step)	0.01 m; 0.001 rad
Sensor range r_s; measurement rate	10 m; 10 Hz
CBF rate γ; risk level ε	0.2; 0.05
Ensemble size N; relaxation α	20; 0.3
Station dwell T_d; planned cycle T_c	8 s; 120 s
Andon delay T_a; response time T_r	15 s; 45 s
Slack penalty λ; feasibility penalty λ_f	10^3; 10^9

5.2. Localisation Scenarios

Three reflector layouts define the localisation quality (Table 3). In scenario A, reflectors line both walls of the aisle. In scenario B, the reflectors along the upper straight and the adjoining parts of the ends are occluded, as by a high-bay storage section, so that the robot dead-reckons for about 26 m in every cycle. In scenario C, two thirds of the reflectors are missing along the whole route and the measurement noise is about three times larger. The fixed margin $\bar{\epsilon}$ is set per scenario at commissioning, close to the 99th percentile of the EKF position error on the obstacle-free route; Table 3 reports that percentile for a drive of twelve runs. Experiment 3 varies the margin, so the comparison does not depend on this choice.

Table 3: Localisation scenarios. Position error: mean and 99th percentile (P99) of the EKF position error in an obstacle-free drive of twelve runs; $\bar{\epsilon}$ is the fixed margin used in Experiment 2.

Scenario	σ_r (m) / σ_b (rad)	Mean (m)	P99 (m)	$\bar{\epsilon}$ (m)
A: full coverage	0.05 / 0.02	0.02	0.05	0.05
B: occluded section	0.05 / 0.02	0.08	0.55	0.60
C: sparse coverage	0.15 / 0.05	0.06	0.26	0.22

5.3. Experiments

Experiment 1 (passing width) tests Corollary 1. The robot drives along the lower straight past one pallet with protrusion δ between 1.0 and 1.9 m. The measurement noise of scenario A is scaled by factors between 1 and 120 to vary the lateral standard deviation. A run counts as a pass if the robot clears the pallet by 6 m within 60 s. We run CC-EKF and ET-CBF with six seeds per condition.

Experiment 2 (shift comparison) runs the six filters of Section 4.6 in the three scenarios. A run consists of three cycles, and each configuration is run 100 times.

Experiment 3 (safety–flow frontier) varies the conservatism of each family of filters in scenario B: the fixed margin $\bar{\epsilon}$ from 0.1 to 0.6 m, the scale c of the adaptive margin from 1.0 to 3.03, and the risk level ε of CC-EKF and ET-CBF from 0.2 to 0.01. Each setting is run 60 times on the first 60 seeds of Experiment 2.

Experiment 4 (ensemble size) runs ET-CBF in scenario B with N between 4 and 40, with the Gaussian gain and with the finite-ensemble gain (10), 60 times per setting.

Experiment 5 (route map) uses the runs of Experiment 2 to map the localisation uncertainty and the protective stops along the route.

Experiment 6 (operator response and recovery) repeats Experiment 2 in scenario B for the four uncertainty-aware filters with a response time T_r of 15 s and of 120 s instead of 45 s, and with the re-localise-first recovery described in Section 4.4, 60 times per setting.

5.4. Statistical Procedures

Runs with the same index share the pallet layout, the operator arrival times and the noise seeds across filters. Rates per hour are totals divided by the simulated operating time. Confidence intervals are 95% percentile intervals from 1000 bootstrap resamples of the runs. All programs are solved with Clarabel [58]. The derivations of Section 4 and their implementation in the simulator were verified symbolically with SymPy [59] and by sampling; Appendix A lists the checks.

6. Results

6.1. Passing Width

Fig. 2 shows the outcome of Experiment 1. Each marker is one combination of measurement noise and pallet protrusion, placed at the lateral standard deviation that the estimator reported while it passed the pallet. The curve is the threshold (14) evaluated with the gain of each filter. At the nominal noise level, the lateral standard deviation is about 2 cm, and both filters pass protrusions up to 1.8 m, against a predicted threshold of 1.81 to 1.83 m; a pallet that leaves 20 cm of the 2 m corridor still lets the robot through. As the standard deviation grows to about 0.2 m, the largest passable protrusion falls to 1.2 to 1.3 m, as predicted. The condition separates passed and blocked runs correctly in 92% of the CC-EKF runs and 90% of the ET-CBF runs, and all disagreements lie within 0.16 m of the predicted threshold. Most disagreements are passes of a protrusion slightly above $\delta_{\max}$, so the first-order condition is necessary only approximately: the standard deviation varies during the pass and the ensemble mean does not sit exactly on the

narrowest cross-section. The condition is therefore best used as a screening rule at commissioning. A protrusion well above $\delta_{\max}$ for the localisation quality of a section will block the vehicle there, while a protrusion below it can be passed but is not guaranteed to be.

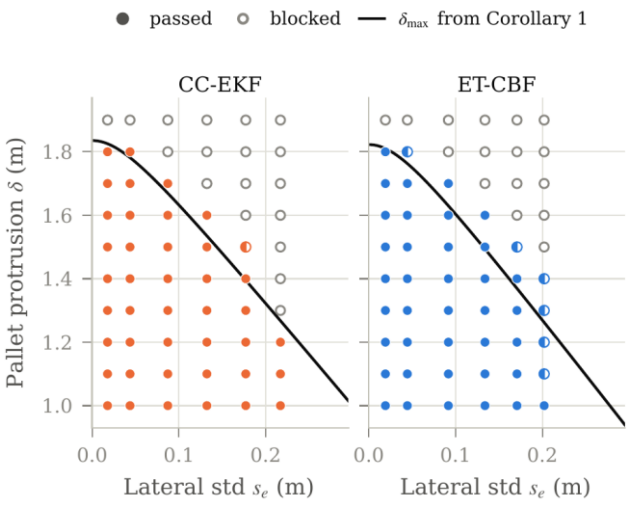


Figure 2: Experiment 1: passing a single pallet in a straight aisle. Filled markers: all runs passed; open markers: all runs ended in a protective stop; half-filled markers: mixed outcome. Line: necessary passing width of Corollary 1.

6.2. Shift Comparison

Table 4 summarises Experiment 2. In scenario A, where the robot is well localised along the whole route, all filters run at the free-flow cycle time, complete every cycle within the planned cycle time and make about 126 station deliveries per hour. The difference between the filters is in safety. The certainty-equivalent filter ignores the localisation error and the odometry noise; its constraints are violated at 36.6% of tight steps, and it infringes the separation 11.6 times per hour. The fixed and adaptive margins reduce the infringements to 0.21 and 0.32 per hour, but their one-step risk stays at 10 to 12%, because their margins cover the localisation error and not the odometry noise of the next step. The two chance-constrained filters, whose margins include both (Proposition 2), cause no infringement, and their realised risk of 2.5% (CC-EKF) and 2.0% (ET-CBF) is below the design level of 5%.

Table 4: Experiment 2: lean performance and safety over 100 runs of three cycles per filter and scenario (about 10 to 14 simulated hours per filter and scenario). Cycle times in seconds; adherence: share of planned cycles completed within $T_c = 120$ s; deliveries, andon calls, infringements, contacts and lane departures per hour of operation; protective stop as a share of travel time; risk: share of tight constraints violated by the true state (per barrier); joint: share of certified steps at which any barrier is violated.

		Flow						Safety				
Scen.	Filter	Cycle mean	Cycle P95	Adher- ence (%)	Deliv- eries/h	Prot. stop (%)	Andon /h	Risk (%)	Joint (%)	Infringe- ments/h	Con- tacts/h	Lane dep./h
A	Oracle (true state)	114.1	115.5	100	126.2	0.0	0	2.7	0.12	0	0	0
	Certainty equivalent	113.9	115.0	100	126.4	0.0	0	36.6	1.43	11.6	0	0
	Fixed margin	114.0	115.3	100	126.3	0.0	0	10.0	0.44	0.21	0	0
	Adaptive margin	114.0	115.3	100	126.3	0.0	0	12.1	0.52	0.32	0	0
	CC-EKF	114.2	115.5	100	126.1	0.0	0	2.5	0.12	0	0	0
	ET-CBF	114.2	115.6	100	126.1	0.0	0	2.0	0.10	0	0	0
B	Oracle (true state)	114.4	116.1	99	125.9	0.0	0.10	2.6	0.13	0	0	0
	Certainty equivalent	114.1	115.0	100	126.2	0.2	0.11	42.7	2.00	23.0	9.3	0.63
	Fixed margin	169.2	361.2	66	85.1	19.4	17.2	0.7	0.45	0.35	0.21	0.64
	Adaptive margin	133.7	235.0	75	107.7	18.2	8.1	3.3	0.30	0.18	0	0.45
	CC-EKF	125.2	176.6	82	115.0	11.5	5.1	3.7	0.28	1.4	0.48	0.86
	ET-CBF	127.8	181.3	81	112.6	13.9	5.8	3.8	0.45	2.6	0.84	2.0
C	Oracle (true state)	114.1	115.7	100	126.2	0.0	0	2.5	0.12	0	0	0
	Certainty	113.9	115.3	100	126.4	0.0	0	39.0	1.60	18.4	1.4	0

equivalent											
Fixed margin	114.6	116.6	100	125.6	0.0	0	0.2	0.02	0.21	0	0
Adaptive margin	115.0	116.1	98	125.2	0.8	0.42	2.5	0.14	0.21	0	0.21
CC-EKF	114.3	116.0	100	126.0	0.0	0	3.0	0.16	0.52	0	0
ET-CBF	114.4	116.0	100	125.9	0.0	0	2.8	0.16	0.42	0	0

Scenario B separates the filters, and it shows that flow and safety have to be read together. The certainty-equivalent filter still runs at free-flow speed, but it infringes the separation 23.0 times per hour, 9.3 of them as contacts, because it steers on an estimate that is up to several decimetres wrong in the occluded section. The fixed margin, sized for the occluded section, applies a 0.6 m stand-off along the whole route. The robot cannot pass pallets that it would pass in scenario A: the mean cycle time rises to 169.2 s, only 66% of the cycles are on time, deliveries fall to 85.1 per hour and the vehicle calls an operator 17.2 times per hour. The adaptive margin applies the stand-off only where the covariance is large, which more than halves the call-outs to 8.1 per hour and raises the adherence to 75%. The chance-constrained filters give the shortest cycles among the implementable uncertainty-aware filters, 125.2 s for CC-EKF and 127.8 s for ET-CBF, with adherence of 82% and 81% and 5.1 and 5.8 andon calls per hour. These filters also cause the most separation events among the uncertainty-aware filters. CC-EKF infringes the separation 1.4 times and leaves the corridor 0.86 times per hour, with 5 contacts in 10.4 simulated hours. ET-CBF infringes 2.6 times and leaves the corridor 2.0 times per hour, with 9 contacts in 10.7 hours. The fixed margin had 3 contacts in 14.1 hours, and the adaptive margin none in 11.1 hours. In paired comparisons over the common seeds, ET-CBF shortens the mean cycle time by 41.3 s (95% interval 35.5 to 47.4 s) and reduces andon calls by 11.3 per hour relative to the fixed margin, but it adds 3.6 separation events per hour (1.5 to 6.0) and 0.6 contacts per hour (0.0 to 1.4). Relative to the adaptive margin, it is 5.9 s faster (3.7 to 8.0 s) and adds 4.0 events per hour (2.1 to 6.4). The filters in Table 4 therefore operate at different points of the trade-off between flow and safety, and the table does not compare them at matched safety; Section 6.4 does.

The joint one-step failure rate on certified steps was 0.12% and 0.10% for CC-EKF and ET-CBF in scenario A and 0.28% and 0.45% in scenario B. It is far below the union bound of about 13%, because most certified steps are far from every barrier. The failures are not spread evenly over the runs. In scenario B, the 5% of runs with the most joint failures account for 49% of them for ET-CBF and 38% for CC-EKF; these are runs in which the estimate drifted further than its reported uncertainty in the occluded section (Section 7).

Between the two chance-constrained filters, CC-EKF is slightly better in scenario B. Its mean cycle time is 2.6 s shorter (0.2 to 4.8 s), and it has 2.3 fewer infringements and lane departures per hour (0.8 to 4.0); the difference in andon calls, 0.7 per hour (-0.4 to 1.8), is not significant. Section 7 examines the cause. In scenario C, where localisation is uniformly mediocre but never lost, all filters except the certainty-equivalent one run within 1 s of free flow; the uncertainty-aware filters infringe the separation between 0.21 and 0.52 times per hour.

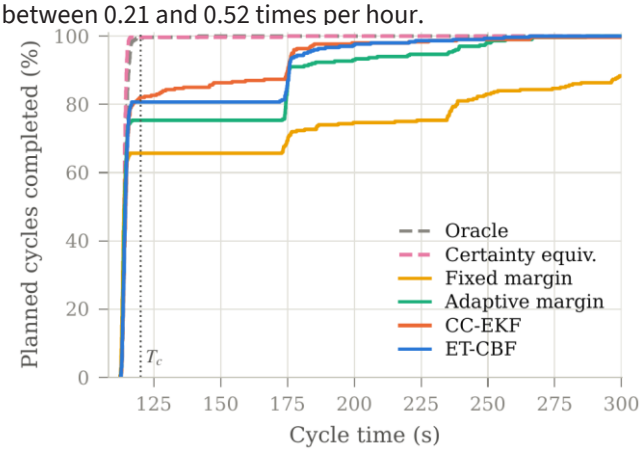


Figure 3: Cumulative distribution of the cycle time in scenario B, as a share of the 300 planned cycles per filter. Cycles that are not completed within the run do not reach 100%. The step near 175 s corresponds to cycles that contain one andon call.

Fig. 3 shows where the cycle time is lost. A cycle either runs close to free flow or contains one or more andon calls, each of which adds the andon delay and the response time of 60 s. The share of cycles without a call is therefore the main quantity that separates the filters, and the mean cycle time follows from it. For a lean operation this is the relevant view: the planned cycle absorbs small slowdowns, but a vehicle that waits a minute for an operator misses its slot.

6.3. Where the Stops Happen

Fig. 4 maps Experiment 5 along the route. In scenario B, the lateral standard deviation of the ETKF is about 0.02 m on most of the loop and rises to 0.32 m at the end of the occluded section; in scenario C it varies between about 0.04 and 0.16 m with the reflector spacing. The lower panel shows the protective-stop time per cycle accumulated along the route in scenario B. For the three uncertainty-aware

filters, between 97% and 99% of the protective-stop time falls in the section between $s = 52$ and 76 m where the lateral standard deviation exceeds 0.1 m. The fixed margin stops wherever a pallet narrows the aisle, and only 6% of its stop time falls in that section; its total stop time is 26.6 s per cycle against 18.5 s for the adaptive margin, 10.7 s for CC-EKF and 13.3 s for ET-CBF.

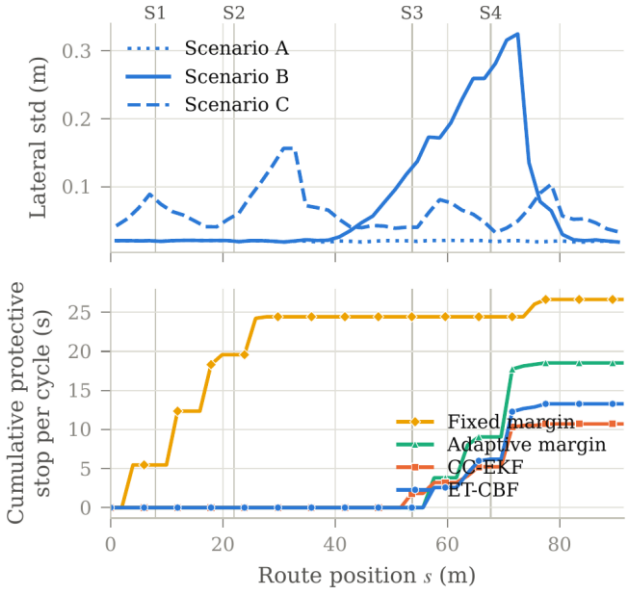


Figure 4: Route map. Top: lateral standard deviation of the ETKF along the route in the three scenarios. Bottom: protective-stop time per cycle, accumulated along the route, in scenario B. Vertical lines: stations S1–S4.

For the uncertainty-aware filters, the stop map therefore points to the cause of the waiting time, the section where localisation is lost, rather than to the locations at which the vehicle stopped. Such a map can direct improvement (kaizen) activities: reflectors added at the end of the occluded section, or a rule that keeps that stretch of aisle clear, would remove most of the waiting. For the fixed margin, the stops are spread over every pallet on the route, which points to 5S discipline everywhere rather than to the real cause.

6.4. Safety–Flow Frontier and Calibration

Fig. 5 shows the trade-off between safety and flow in scenario B, where separation events are infringements and lane departures. Each family of filters traces a curve as its conservatism varies. With $\bar{\epsilon} = 0.1$ m, the fixed margin runs at free flow (114.2 s) but causes 10.0 separation events per hour, 4.38 of them contacts. Bringing the events down to 2.4 per hour requires $\bar{\epsilon} = 0.45$ m, a mean cycle time of 143.8 s and 11.3 andon calls per hour. The three uncertainty-aware filters lie on a common frontier well below this curve. At about the same event rate, CC-EKF with $\varepsilon = 0.05$ has 2.4 events per hour, a mean cycle time of 125.5 s and 4.9 andon calls per hour, and the contact rates of the two settings are also similar (0.42 and 0.48 per hour). Below one event per hour, the adaptive margin with $c = 2.45$ (0.9 events per hour, 134.7 s), CC-EKF with $\varepsilon = 0.01$ (0.7, 133.6 s) and ET-CBF with $\varepsilon = 0.02$ (0.6, 141.6 s) compare with 171.4 s for the fixed margin at 0.5 events per hour. The gain in flow therefore comes from tying the margin to the uncertainty of the estimator. Among the uncertainty-aware filters, the curves overlap within their sampling intervals on this route.

The chance constraints differ from the adaptive margin in the interpretation of the tuning parameter. Panel (c) shows that the realised one-step risk of both chance-constrained filters follows the design level and stays below it. CC-EKF realises 0.8, 1.1, 3.5, 6.6 and 12.5% at design levels of 1, 2, 5, 10 and 20%, and ET-CBF realises 0.1, 0.9, 3.5, 7.1 and 16.4%. The oracle realises 2.6% at $\varepsilon = 0.05$ (Table 4), so part of this conservatism comes from the tangent condition of Lemma 1 and not from the estimators. The scale c of the adaptive margin has no such interpretation. With $c = 2.45$ the margin covers a 95% confidence region of the position, yet the realised one-step risk is 3.1% in scenario B and 12.1% in scenario A, where the margin ignores the odometry noise of the next step; with $c = 1.645$ it is 9.1%. A plant can therefore set ε from an accepted risk per step and expect it to hold across sections and layouts, whereas the same margin setting implies different risks in different plants. At the smallest risk levels, ET-CBF with 20 members is considerably more conservative than CC-EKF, with a mean cycle time of 155.3 s against 133.6 s at $\varepsilon = 0.01$. In the tail, the finite-ensemble gain grows faster than the Gaussian one ($\kappa_{20} = 2.60$ against 2.33 at $\varepsilon = 0.01$), and the fluctuation of the sample variance adds protective stops.

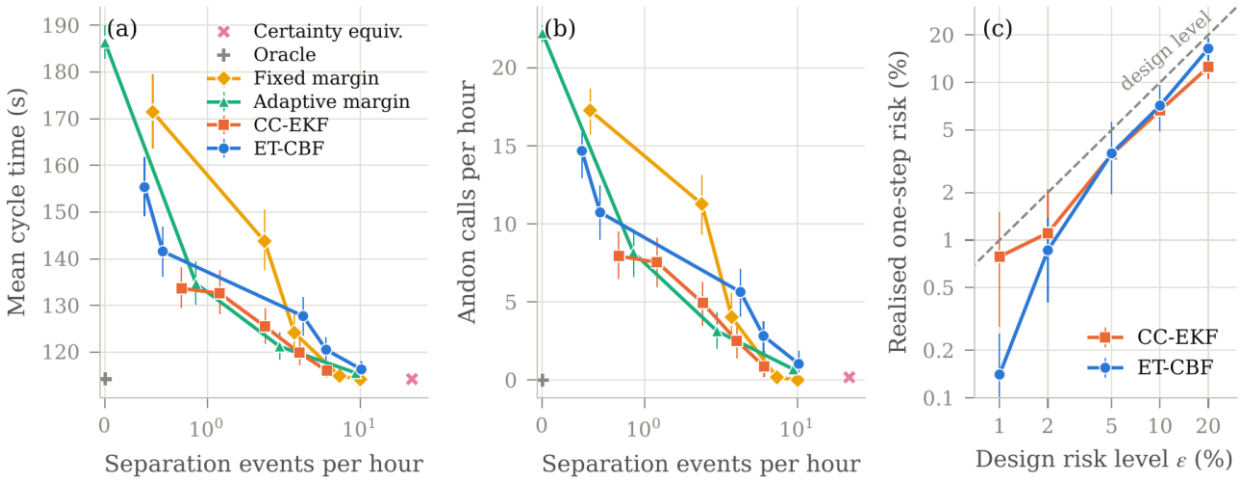


Figure 5: Experiment 3 in scenario B. (a) Mean cycle time and (b) andon calls per hour against the rate of separation infringements and lane departures, for each family of filters as its conservatism varies ($\bar{\epsilon}$ from 0.1 to 0.6 m, c from 1.0 to 3.03, ε from 0.2 to 0.01); certainty equivalent and oracle shown for reference. Horizontal axis linear below one event per hour and logarithmic above. Bars: 95% bootstrap intervals of the vertical quantity. (c) Realised one-step risk of the chance-constrained filters against the design level.

6.5. Ensemble Size

Fig. 6 shows Experiment 4. With the Gaussian gain, small ensembles exceed the design risk: the realised one-step risk is 10.0% with four members and 8.5% with six, and the vehicle causes 10.1 separation events per hour with four members. This matches the idealised Monte Carlo check of Appendix A, in which the Gaussian gain yields 11.9% for $N = 4$. The finite-ensemble gain (10) brings the realised risk to 3.7% and 2.8%, and the events to 4.0 and 3.1 per hour. The price is flow: with four members the gain rises from 1.64 to 2.63, the mean cycle time from 126.2 to 148.5 s and the andon calls from 4.8 to 12.4 per hour. The price falls quickly with the ensemble size. From about 20 members, the Gaussian gain already keeps the risk below the design level on this route (3.9% for $N = 20$ and 3.4% for $N = 40$), and the finite-ensemble gain costs about 2 to 3 s per cycle. The solve time does not depend on the ensemble size and stayed between 0.8 and 0.9 ms, because the QR reduction keeps every cone at dimension five. For deployment, we therefore suggest 20 to 40 members with the finite-ensemble gain; smaller ensembles are safe with that gain but cost flow.

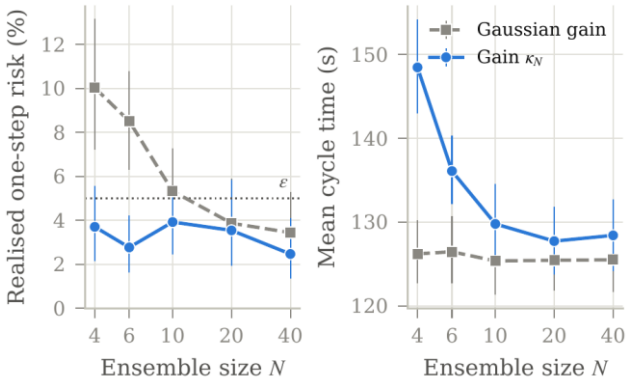


Figure 6: Experiment 4: ET-CBF in scenario B with ensemble sizes from 4 to 40, with the Gaussian gain and with the finite-ensemble gain (10). Bars: 95% bootstrap intervals.

6.6. Operator Response and Recovery

Table 5 shows Experiment 6. The response time sets the cost of each andon call, so it scales the mean cycle time of every filter, but it does not change their order. With $T_r = 15$ s, the mean cycle time of the fixed margin falls from 171.4 to 146.0 s, and with $T_r = 120$ s it rises to 230.6 s; for ET-CBF the corresponding values are 127.7, 122.3 and 143.1 s. Cycle adherence hardly changes with T_r, because one call already pushes a cycle beyond the planned cycle time even with the fastest response ($T_{\text{free}} + T_a + 15\,\text{s} \approx 143\,\text{s}$). The re-localise-first recovery is the more demanding variant. A vehicle whose stop has a physical cause calls at least twice before the operator clears the aisle, so the call rate rises by a factor of 1.6 to 2.2, to 9.7 per hour for ET-CBF

and 27.1 per hour for the fixed margin, and adherence falls to 66% for ET-CBF, 64% for CC-EKF and 35% for the adaptive margin. In all settings, the fixed margin has the longest cycles and the chance-constrained filters the shortest, and the chance-constrained filters keep their higher rate of separation events. The comparison in Table 4 is therefore not an artefact of the response time or the recovery procedure, although the size of the differences depends on both.

Table 5: Experiment 6: sensitivity of the scenario B results to the operator response time T_r and to the recovery procedure, over 60 runs per setting; $T_r = 45\,\mathrm{s}$ is the base case (the first 60 runs of Experiment 2). Events: infringements and lane departures per hour. Re-loc. first: the operator first only re-localises the vehicle (Section 4.4).

Filter	Setting	Cycle mean (s)	Adh. (%)	Andon /h	Events /h	Con-tacts/h
Fixed	$T_r = 45\,\mathrm{s}$	171.4	67	17.3	0.47	0.12
	$T_r = 15\,\mathrm{s}$	146.0	66	20.4	0.82	0.27
	$T_r = 120\,\mathrm{s}$	230.6	67	12.7	0.43	0.26
	Re-loc. first	229.6	63	27.1	0.70	0.26
Adaptive	$T_r = 45\,\mathrm{s}$	134.7	76	8.2	0.89	0
	$T_r = 15\,\mathrm{s}$	125.7	76	8.7	1.1	0
	$T_r = 120\,\mathrm{s}$	157.6	76	7.0	0.63	0
	Re-loc. first	171.9	35	17.8	0.93	0.12
CC-EKF	$T_r = 45\,\mathrm{s}$	125.5	82	4.9	2.4	0.48
	$T_r = 15\,\mathrm{s}$	121.0	83	5.5	2.8	0.17
	$T_r = 120\,\mathrm{s}$	134.9	82	4.7	1.9	0.30
	Re-loc. first	140.6	64	10.4	2.8	0.85
ET-CBF	$T_r = 45\,\mathrm{s}$	127.7	81	5.6	4.2	0.63
	$T_r = 15\,\mathrm{s}$	122.3	81	5.9	4.4	1.0
	$T_r = 120\,\mathrm{s}$	143.1	81	5.2	2.9	0.70
	Re-loc. first	140.3	66	9.7	4.1	1.3

7. Discussion

7.1. Implications for Lean Intralogistics

The results support three practical conclusions. First, the safety layer of a vehicle is part of the cycle-time budget. In scenario B, the choice of filter moves the share of on-time cycles from 66% to 82% without any change to the route, the vehicle or the plant, although the faster filters also accept more separation events. A plant that sizes a fixed protective field for its worst section incurs this cost along the whole route as waiting time at every pallet that is set down slightly outside its marking. At matched safety, any margin that follows the covariance of the estimator removes much of this cost (Section 6.4), and the chance constraint adds a tuning parameter whose meaning carries over from one plant to the next. Second, Corollary 1 gives a quantitative link between two lean disciplines that are usually managed separately: the localisation infrastructure and the 5S state of the aisle. The upper limit on the tolerated protrusion falls by 2κ times the lateral standard deviation, 3.3 to 3.5 times at $\varepsilon = 0.05$ depending on the gain, so each section of a route can be given an explicit screening limit for floor markings. Third, the protective-stop rule of Section 4.4 gives the safety filter a role analogous to jidoka. The vehicle stops when it cannot certify a safe command rather than when it has already come too close, the stop is logged with its position, and the resulting map points to the root cause.

7.2. How Far the Assumptions Hold, and ET-CBF Against CC-EKF

In this study, the ensemble filter does not outperform the linearised filter, and in the scenario with an occluded section it is slightly worse. We trace this to the calibration of the estimators in the dead-reckoning section, which is also where Assumption 1 is weakest. In a separate run of 30 obstacle-free cycles per estimator in scenario B, we compared the mean squared position error with the mean variance reported by the estimator in each section of the route. On the well-localised part of the route, the ratio was 0.99 to 1.07 for the EKF and 0.88 to 0.96 for the ETKF. In the occluded section it was 0.70 to 0.83 for the EKF, which therefore overstates its uncertainty there, and 0.85 to 1.15 for the ETKF, which understates it towards the end of the section. An overstated variance makes CC-EKF more cautious exactly where the risk is, whereas an understated variance lets ET-CBF approach pallets and boundaries with a little less margin than the risk level requires. With 20 members, the sample variance also fluctuates from step to step, which may explain the higher protective-stop share of ET-CBF (13.9% against 11.5%). The concentration of joint failures in a few runs (Section 6) points the same way: Assumption 1 holds well on average but fails in episodes in which the ensemble is overconfident, and these episodes produce most of the separation events. The per-step risk statement of Proposition 1 should therefore be read as a design target that holds on average, not as a guarantee for every step.

The advantage of the ensemble formulation lies elsewhere. It needs no Jacobians of the barriers or of the dynamics, so new barrier types, such as the outline of a

pallet or a map-based keep-out zone, can be added by evaluating them on the members, and it extends to estimators that are already ensemble-based. Its computational cost is comparable: the mean solve time was 0.64 ms per step for ET-CBF and 0.65 ms for CC-EKF, against 0.27 ms for the certainty-equivalent filter, well within the 100 ms control period. The present plant model is only mildly nonlinear, which is the regime in which the EKF is consistent. Settings with strongly nonlinear barriers or multimodal localisation, for example after a kidnapping or in symmetric aisles, are where the ensemble can be expected to help, and they remain to be tested.

7.3. Limitations

The study has several limitations. First, the evidence is from simulation. The robot is a kinematic model with Gaussian odometry noise, the reflector map is exact, and systematic errors such as wheel slip, map distortion or a miscalibrated scanner are not represented; such errors would bias the estimate in a way that neither filter covariance captures. Second, operators are tracked without error by the plant infrastructure and do not react to the vehicle, and pallets are discs of known position; uncertainty in the obstacles would add to the variance in (8) and could be included in the same way as the process noise. Third, the risk statement is a per-barrier, one-step statement under Assumption 1. It is not a collision probability: localisation errors persist over many steps, so the rate of infringements per hour is not bounded by ε and has to be set through the frontier of Fig. 5, and a bound on the risk over a cycle would require the trajectory-level analysis of [49, 50]. Fourth, safety events are rare in absolute terms. Each filter accumulates only 10 to 14 simulated hours per scenario in Experiment 2, so the contact rates rest on a few events each and have wide intervals; a safety case for a real deployment would need far more exposure. Fifth, we study one vehicle. With several vehicles on the route, a protective stop delays the followers, and the interaction between vehicles would have to be priced or coordinated [9, 60]. Sixth, the andon model is stylised. Experiment 6 shows how the results depend on the response time and the recovery procedure, but the response time is deterministic within each setting and the operator always succeeds. Finally, the model of this study is only mildly nonlinear, which is the regime in which an EKF is consistent; we therefore cannot claim an accuracy advantage of the ensemble filter here, and the comparison with CC-EKF should be read as a test of equivalence.

8. Conclusion

This paper integrated uncertainty-aware safety filtering with the evaluation of a lean milk run and reported flow and safety side by side. The study leads to three main findings. First, the safety layer of the vehicle takes a large share of the cycle-time budget. In the scenario with an occluded section, a fixed margin sized for the worst section of the route gave 41.3 s longer cycles and 11.3 more operator calls per hour than the ensemble filter, while the ensemble filter incurred 3.6 more separation events per hour. Second, when the filters are compared at matched safety, the gain in flow comes from tying the margin to the uncertainty of the estimator. The adaptive margin and the two chance constraints share one frontier well below that of the fixed margin; only the chance constraints have a tuning parameter whose realised one-step risk follows its design value, and the ensemble filter performed close to, though slightly behind, the linearised chance constraint in this mildly nonlinear setting. Third, the protective-field view yields practical aids for commissioning. The necessary passing width of Corollary 1 screens which pallet protrusions a section of the route can tolerate for a given localisation quality, and the route map of protective stops directs improvement work to the section where localisation is lost rather than to the locations at which the vehicle stopped.

The risk statements of the method are per-barrier, one-step statements under a Gaussian consistency assumption that holds on average but fails in episodes of overconfidence; they are design targets, not guarantees of collision-free operation. Further work should test the filter on a physical vehicle with real localisation errors and much longer exposure, extend it to several vehicles that share a route, bound the risk over a whole cycle, and examine strongly nonlinear barriers and multimodal localisation, where the ensemble formulation may have an advantage over linearisation.

CRediT Authorship Contribution Statement

Toheeb Babatunde Ibrahim: Conceptualization, Methodology, Software, Writing – original draft. **Olasubomi Akanbi:** Formal analysis, Validation, Writing – review & editing. **Joseph Oladejo:** Visualization. **Warith Akinade:** Software.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data Availability

The simulation code, the experiment scripts, the raw results of all runs, the analysis script that produces every number, table and figure in this paper, and the verification script are provided in the supplementary archive of this manuscript. A public repository will be linked on publication.

Declaration of Generative AI and AI-Assisted Technologies

During the preparation of this work the authors used Claude (Anthropic) to assist with the literature search, the simulation code, the drafting of the text and the checking of the text for grammatical errors and mistakes. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Appendix A. Symbolic and Numerical Verification

We verified the derivations of Section 4, and their implementation in the simulator, with the script `sympy_validate.py`, which accompanies the simulation code. The script uses SymPy [59] for symbolic differentiation, expansion, eigenvalues, probability distributions and equation solving, and NumPy and SciPy for sampling. Table A.1 lists the 27 checks; all of them pass. Three kinds of check are used. Symbolic checks (marked *exact*) reduce the difference between a derived expression and the expression in the paper to zero. Implementation checks evaluate the symbolic expressions, compiled with `lambdify`, at random states and compare them with the functions of the simulator. Inequalities that hold for all states, such as the tangent conditions, are checked by sampling and reported as the smallest slack; a nonnegative slack, up to rounding, means that no counterexample was found.

The checks also give the constants quoted in the text. The largest displacement of the reference point in one step is $\Delta t\sqrt{v_{\max}^2 + \ell^2\omega_{\max}^2} = 0.162$ m, so the residual of the outer boundary in Lemma 1 is at most 3.26 mm per step at the input bounds and 1.80 mm at cruise speed; the largest sampled residual was 3.24 mm. The finite-ensemble gain (10), computed from the symbolic Student-t distribution function, is 2.631, 2.177, 1.923, 1.772 and 1.706 for $N = 4$, 6, 10, 20 and 40 at $\varepsilon = 0.05$, against the Gaussian value of 1.645. For independent Gaussian draws, which ETKF members are not, the Gaussian gain gives a realised risk of 11.9%, 7.5% and 6.1% for $N = 4$, 10 and 20, whereas κ_N gives 4.96%, 4.98% and 4.90% (Monte Carlo, 4×10^5 draws each). In the limit of a large lateral standard deviation, the derivative of $\delta_{\max}$ in (14) with respect to s_e tends to -2κ. The Monte Carlo checks of the union bound (C17) and of the gain (C16) agree with the theory within their sampling error of about 10^{-3}.

Table A.1: Symbolic and numerical verification of the method and the simulator. Exact: the difference between the derived and the stated expression simplifies to zero in SymPy. Max. error: largest absolute difference between the symbolic expression and the simulator. Min. slack: smallest value of an inequality over the samples (negative values within rounding or Monte Carlo error).

Check	Result in paper	Statement verified	Method	Outcome
C1	Eq. (1)	Reference-point kinematics yield $G(x)$	symbolic differentiation	exact
C2	Eq. (1)	Simulator G_q equals symbolic G_q	lambdified vs. simulator	max. error 3.5×10^{-18}
C3	Eq. (4)	Disc barrier has unit gradient	symbolic	exact
C4	Lemma 1	Disc-barrier Hessian has eigenvalues $\{0, 1/d\}$, so the barrier is convex	symbolic eigenvalues	exact
C5	Lemma 1	Tangent inequality $h(q+\Delta) \geq h(q) + n^{\mathsf{T}}\Delta$ for discs	random sampling	min. slack 2.0×10^{-13}
C6	Lemma 1	Tangent inequality for the inner aisle boundary	random sampling	min. slack -1.0×10^{-15}

Check	Result in paper	Statement verified	Method	Outcome
C7	Lemma 1	Outer-boundary residual $\leq \| \Delta q \|^2/\big(2(R-w)\big)$	random sampling	min. slack 4.0×10^{-9}
C8	Eqs. (6)–(7)	$\psi = a_i + b_i^\top u + n^\top w$	symbolic expansion	exact
C9	Eq. (7)	Simulator a_i, b_i equal symbolic values	lambdified vs. simulator	max. error 8.9×10^{-16}
C10	Eq. (8)	Sample mean of ψ_i equals $\bar{a} + \bar{b}^\top u$	symbolic, $N = 4$	exact
C11	Eq. (8)	Sample variance equals $\| a' + B'u \|^2/(N-1)$	symbolic, $N = 4$	exact
C12	Eq. (8)	$n^\top(\sigma_w^2 I)n = \sigma_w^2$ for unit n	symbolic	exact
C13	Eq. (9)	QR-reduced cone reproduces the sample mean and standard deviation	random ensembles, $N = 4$–40	max. error 1.7×10^{-16}
C14	Prop. 1	$\mathbb{P}[\psi < 0] = \varepsilon$ when $\mu = \Phi^{-1}(1-\varepsilon)\sigma$	SymPy Gaussian cdf	max. error 0
C15	Eq. (10)	κ_N from the symbolic Student-t cdf equals the simulator gain	SymPy root finding	max. error 8.9×10^{-16}
C16	Eq. (10)	κ_N gives risk 0.05 for independent Gaussian draws (not ETKF members)	Monte Carlo, $N = 4{,}10{,}20$	max. error 9.9×10^{-4}
C17	Prop. 1	Joint risk $\leq \sum_j \varepsilon_j$ for correlated barriers	Monte Carlo, 2–4 barriers	min. slack -6.9×10^{-4}
C18	Eq. (12)	$\hat{\mu} - \kappa\hat{\sigma} = \gamma\big(\bar{h} - m(u)\big) + \bar{b}^\top u$	symbolic	exact
C19	Eq. (13)	$m^0 = \kappa\sqrt{s_h^2 + \sigma_w^2/\gamma^2}$	symbolic	exact
C20	Eq. (13)	$\sigma_w \leq \hat{\sigma}(u) \leq \hat{\sigma}(0) + \| B'u \|/\sqrt{N-1}$	random ensembles	min. slack 2.2×10^{-5}
C21	Prop. 2	First-order barrier variance $s_h^2 = n^\top P_q n$	symbolic	exact
C22	Eq. (14)	Summing the opposed constraints yields $\delta_{\max}$	symbolic solution	exact
C23	Cor. 1	Largest passed protrusion lies within 0.16 m of $\delta_{\max}$	SymPy $\delta_{\max}$ vs. Exp. 1	max. error 0.131
C24	Eq. (15)	Relaxed spread equals $(1-\alpha)s^a + \alpha s^b \leq s^b$	symbolic	exact
C25	Sec. 4.5	Variance growth $\rho^{fT_d} \approx 49.6$ over one dwell	SymPy evaluation	max. error 1.4×10^{-3}
C26	Eq. (9)	Sign of the simulator cone residual equals sign of $\hat{\mu} - \kappa\hat{\sigma}$	random ensembles	max. error 0
C27	Sec. 4.4	Tracking cost $\leq v_{\max}^2 + 4\ell^2\omega_{\max}^2 < 4$, so $\lambda_f = 10^9$ yields a zero-slack command when one exists	SymPy, vertices of the input box	max. error 0

REFERENCES

[1] J. P. Womack, D. T. Jones, D. Roos, The Machine That Changed the World, Rawson Associates, New York, 1990.

[2] T. Ohno, Toyota Production System: Beyond Large-Scale Production, Productivity Press, Portland, OR, 1988.

[3] J. K. Liker, The Toyota Way: 14 Management Principles from the World's Greatest Manufacturer, McGraw-Hill, New York, 2004.

[4] S. Emde, N. Boysen, Optimally routing and scheduling tow trains for JIT-supply of mixed-model assembly lines, European Journal of Operational Research 217 (2) (2012) 287–299.

[5] D. Battini, N. Boysen, S. Emde, Just-in-time supermarkets for part supply in the automobile industry, Journal of Management Control 24 (2) (2013) 209–217.

[6] N. Boysen, S. Emde, M. Hoeck, M. Kauderer, Part logistics in the automotive industry: Decision problems, literature review and research agenda, European Journal of Operational Research 242 (1) (2015) 107–120.

[7] H. S. Kilic, M. B. Durmusoglu, Advances in assembly line parts feeding policies: A literature review, Assembly Automation 35 (1) (2015) 57–68.

[8] E. Klenk, S. Galka, W. A. Günthner, Operating strategies for in-plant milk-run systems, IFAC-PapersOnLine 48 (3) (2015) 1882–1887.

[9] G. Fragapane, R. De Koster, F. Sgarbossa, J. O. Strandhagen, Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda, European Journal of Operational Research 294 (2) (2021) 405–426.

[10] M. De Ryck, M. Versteyhe, F. Debrouwere, Automated guided vehicle systems, state-of-the-art control algorithms and

techniques, Journal of Manufacturing Systems 54 (2020) 152–173.

[11] E. Hofmann, M. Rüsch, Industry 4.0 and the current status as well as future prospects on logistics, Computers in Industry 89 (2017) 23–34.

[12] S.-V. Buer, J. O. Strandhagen, F. T. S. Chan, The link between Industry 4.0 and lean manufacturing: Mapping current research and establishing a research agenda, International Journal of Production Research 56 (8) (2018) 2924–2940.

[13] A. Sanders, C. Elangeswaran, J. Wulfsberg, Industry 4.0 implies lean manufacturing: Research activities in Industry 4.0 function as enablers for lean manufacturing, Journal of Industrial Engineering and Management 9 (3) (2016) 811–833.

[14] D. Kolberg, D. Zühlke, Lean automation enabled by Industry 4.0 technologies, IFAC-PapersOnLine 48 (3) (2015) 1870–1875.

[15] International Organization for Standardization, ISO 3691-4:2023 industrial trucks — safety requirements and verification — part 4: Driverless industrial trucks and their systems, ISO, Geneva (2023).

[16] A. A. Tubis, H. Poturaj, Risk related to AGV systems—open-access literature review, Energies 15 (23) (2022) 8910.

[17] S. Thrun, W. Burgard, D. Fox, Probabilistic Robotics, MIT Press, Cambridge, MA, 2005.

[18] A. D. Ames, X. Xu, J. W. Grizzle, P. Tabuada, Control barrier function based quadratic programs for safety critical systems, IEEE Transactions on Automatic Control 62 (8) (2017) 3861–3876.

[19] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, P. Tabuada, Control barrier functions: Theory and applications, in: Proceedings of the 18th European Control Conference (ECC), 2019, pp. 3420–3431.

[20] K. P. Wabersich, A. J. Taylor, J. J. Choi, K. Sreenath, C. J. Tomlin, A. D. Ames, M. N. Zeilinger, Data-driven safety filters: Hamilton-Jacobi reachability, control barrier functions, and predictive methods for uncertain systems, IEEE Control Systems Magazine 43 (5) (2023) 137–177.

[21] K.-C. Hsu, H. Hu, J. F. Fisac, The safety filter: A unified view of safety-critical control in autonomous systems, Annual Review of Control, Robotics, and Autonomous Systems 7 (2024).

[22] F. Ferraguti, M. Bertuletti, C. T. Landi, M. Bonfè, C. Fantuzzi, C. Secchi, A control barrier function approach for maximizing performance while fulfilling to ISO/TS 15066 regulations, IEEE Robotics and Automation Letters 5 (4) (2020) 5921–5928.

[23] F. Parma, C. Tonola, N. Pedrocchi, M. Beschi, Embedding ISO 10218 safety compliance in robots via control barrier functions for human-robot collaboration, arXiv:2606.13203 (2026).

[24] Z. Zeng, S. Chen, X. Kong, X. Li, C. Zhang, G. Yang, Revised control barrier function with sensing of threats from relative velocity between humans and mobile robots, Sensors 25 (13) (2025) 4005.

[25] S. Farrell, C. Li, H. Yu, R. Yoshimitsu, S. Gao, H. I. Christensen, Safe human robot navigation in warehouse scenario, arXiv preprint arXiv:2503.21141 (2025).

[26] A. Miksits, F. S. Barbosa, M. Lindhé, J. Araújo, K. H. Johansson, Safe navigation of networked robots under localization uncertainty using robust control barrier functions, in: Proceedings of the 62nd IEEE Conference on Decision and Control (CDC), 2023.

[27] S. Dean, A. J. Taylor, R. K. Cosner, B. Recht, A. D. Ames, Guaranteeing safety of learned perception modules via measurement-robust control barrier functions, in: Proceedings of the 2020 Conference on Robot Learning (CoRL), PMLR 155, 2021, pp. 654–670.

[28] R. K. Cosner, A. W. Singletary, A. J. Taylor, T. G. Molnar, K. L. Bouman, A. D. Ames, Measurement-robust control barrier functions: Certainty in safety with uncertainty in state, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2021, pp. 6286–6291.

[29] E. Daş, R. Nanayakkara, X. Tan, R. M. Bena, J. W. Burdick, P. Tabuada, A. D. Ames, Safe navigation under state uncertainty: Online adaptation for robust control barrier functions, IEEE Robotics and Automation Letters 11 (3) (2026) 3398–3405. doi:10.1109/LRA.2026.3653366.

[30] A. Clark, Control barrier functions for stochastic systems, Automatica 130 (2021) 109688.

[31] S. Yaghoubi, G. Fainekos, T. Yamaguchi, D. Prokhorov, B. Hoxha, Risk-bounded control with Kalman filtering and stochastic barrier functions, in: Proceedings of the 60th IEEE Conference on Decision and Control (CDC), 2021, pp. 5213–5219.

[32] M. Vahs, C. Pek, J. Tumova, Belief control barrier functions for risk-aware control, IEEE Robotics and Automation Letters 8 (12) (2023) 8565–8572.

[33] M. Vahs, J. Tumova, Risk-aware control for robots with non-Gaussian belief spaces, in: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2024, pp. 11661–11667.

[34] M. Kishida, Risk-aware control of discrete-time stochastic systems: Integrating Kalman filter and worst-case CVaR in control barrier functions, arXiv:2312.15638 (2023).

[35] T. B. Ibrahim, O. Akanbi, J. Oladejo, W. Akinade, Control barrier functions coupled with an ensemble transform Kalman filter for calibrated chance-constrained safety filtering, Manuscript in preparation (2026).

[36] C. H. Bishop, B. J. Etherton, S. J. Majumdar, Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects, Monthly Weather Review 129 (3) (2001) 420–436.

[37] B. R. Hunt, E. J. Kostelich, I. Szunyogh, Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter, Physica D: Nonlinear Phenomena 230 (1–2) (2007) 112–126.

[38] W. Luo, W. Sun, A. Kapoor, Multi-robot collision avoidance under uncertainty with probabilistic safety barrier certificates, in: Advances in Neural Information Processing Systems 33 (NeurIPS 2020), 2020.

[39] P. R. Wurman, R. D'Andrea, M. Mountz, Coordinating hundreds of cooperative, autonomous vehicles in warehouses, AI Magazine 29 (1) (2008) 9–19.

[40] K. Azadeh, R. De Koster, D. Roy, Robotized and automated warehouse systems: Review and recent developments, Transportation Science 53 (4) (2019) 917–945.
[41] N. Boysen, R. De Koster, F. Weidinger, Warehousing in the e-commerce era: A survey, European Journal of Operational Research 277 (2) (2019) 396–411.
[42] S. Shingo, Zero Quality Control: Source Inspection and the Poka-Yoke System, Productivity Press, Portland, OR, 1986.
[43] S. Molina, A. Mannucci, M. Magnusson, D. Adolfsson, H. Andreasson, M. Hamad, S. Abdolshah, R. T. Chadalavada, L. Palmieri, T. Linder, et al., The ILIAD safety stack: Human-aware infrastructure-free navigation of industrial mobile robots, IEEE Robotics & Automation Magazine 31 (3) (2024) 48–59.
[44] E. Brorsson, K. Ceder, Z. Zhang, S. F. Roselli, E. Erős, M. Dahl, B. Alenljung, J. Lindblom, T. Bui, E. Dean, L. Svensson, K. Bengtsson, P.-L. Götvall, K. Åkesson, Infrastructure-based autonomous mobile robots for internal logistics: Challenges and future perspectives, arXiv:2512.15215 (2025).
[45] ISO, ISO/TS 15066:2016 robots and robotic devices — collaborative robots, International Organization for Standardization, Geneva (2016).
[46] L. Blackmore, M. Ono, B. C. Williams, Chance-constrained optimal path planning with obstacles, IEEE Transactions on Robotics 27 (6) (2011) 1080–1094.
[47] J. Zhao, Z. Cai, X. Yin, Output-feedback safe control of discrete-time stochastic systems with chance constraints, arXiv:2604.12956 (2026).
[48] H. Parwana, D. Panagou, FORESEE: Prediction with expansion-compression unscented transform for online policy optimization, arXiv:2209.12644 (2022).
[49] R. K. Cosner, P. Culbertson, A. J. Taylor, A. D. Ames, Robust safety under stochastic uncertainty with discrete-time control barrier functions, in: Robotics: Science and Systems XIX, 2023.
[50] R. K. Cosner, P. Culbertson, A. D. Ames, Bounding stochastic safety: Leveraging Freedman's inequality with discrete-time control barrier functions, IEEE Control Systems Letters 8 (2024) 1937–1942.
[51] M. F. Reis, A. P. Aguiar, P. Tabuada, Control barrier function-based quadratic programs introduce undesirable asymptotically stable equilibria, IEEE Control Systems Letters 5 (2) (2021) 731–736.
[52] J. Grover, C. Liu, K. Sycara, The before, during, and after of multi-robot deadlock, The International Journal of Robotics Research 42 (6) (2023) 317–336.
[53] G. Evensen, Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, Journal of Geophysical Research: Oceans 99 (C5) (1994) 10143–10162.
[54] M. K. Tippett, J. L. Anderson, C. H. Bishop, T. M. Hamill, J. S. Whitaker, Ensemble square root filters, Monthly Weather Review 131 (7) (2003) 1485–1490.
[55] P. Sakov, P. R. Oke, Implications of the form of the ensemble transformation in the ensemble square root filters, Monthly Weather Review 136 (3) (2008) 1042–1053.
[56] J. S. Whitaker, T. M. Hamill, Evaluating methods to account for system errors in ensemble data assimilation, Monthly Weather Review 140 (9) (2012) 3078–3089.
[57] A. Agrawal, K. Sreenath, Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation, in: Robotics: Science and Systems XIII, 2017.
[58] P. J. Goulart, Y. Chen, Clarabel: An interior-point solver for conic programs with quadratic objectives, arXiv:2405.12762 (2024).
[59] A. Meurer, C. P. Smith, M. Paprocki, et al., SymPy: Symbolic computing in Python, PeerJ Computer Science 3 (2017) e103.
[60] L. Wang, A. D. Ames, M. Egerstedt, Safety barrier certificates for collisions-free multirobot systems, IEEE Transactions on Robotics 33 (3) (2017) 661–674.