SENSOR INFORMATION PROCESSING AND ADAPTIVE CONTROL FOR REAL-TIME MOTION AUTOMATION OF AUTONOMOUS ROBOTS: MULTI-RATE STATE ESTIMATION WITH DELAYED MEASUREMENTS

Sherzod A. Yarashev
Tashkent University
of Information Technologies named after
Muhammad al-Khwarizmi, Tashkent, Uzbekistan
ORCID: 0009-0009-5053-942X
Jurabek M. Abdiev
Arifov Institute of Ion-Plasma and Laser
Technologies, Academy of Sciences of the
Republic of Uzbekistan, Tashkent, Uzbekistan
E-mail: fiztexabdiev@gmail.com
ORCID: 0000-0003-4524-3439
https://doi.org/10.5281/zenodo.23906570

Abstract. Autonomous mobile robots must turn fast, drifting proprioceptive data and slow, delayed, outlier-prone pose fixes into a state estimate that a controller can use while payload and friction change. This simulation study asks what delayed-measurement replay, innovation gating and a projection-based adaptive term add to a baseline PID loop for a differential-drive robot. A multi-rate extended Kalman filter (EKF) fuses 100 Hz encoders and gyroscope with 10 Hz pose fixes (60 ms latency, 2 % outliers); boundedness of the adaptive term is analysed for a scalar channel. A figure-eight path with a payload step was simulated with 40 paired seeds per condition, with ablations, an obstacle scenario and timing of the Python code. In simulation, the full pipeline gave a position RMSE of 0.024 m, against 0.372 m for encoders and gyroscope only; omitting delay handling or gating raised it by 45 % and 57 %. All 499 outlier fixes were rejected, with 1.3 % of valid fixes, and the filter was overconfident (mean NEES 10.2, expected 3). The adaptive term shortened speed recovery after the step from 1.02 s to 0.68 s but left the tracking error unchanged (0.038 m) and raised actuator saturation from 2.4 % to 6.3 %. The safety filter prevented contact under the full pipeline (0 of 40 runs) but not with encoder-gyroscope estimates (11 of 40). Compute time averaged 0.4 ms per cycle. The plant model is simple, the boundedness argument does not cover the implemented system, and hardware validation is pending.

Keywords: state estimation; extended Kalman filter; delayed measurements; adaptive control; sensor fusion; autonomous mobile robots; real-time systems

1. Introduction

Autonomous mobile robots are deployed in manufacturing, logistics, agriculture and inspection, where automation and control technology is a main lever for efficiency and reliability in industrial enterprises. Task success depends on how quickly sensor data reach the estimator, how reliable the state estimate is, and how little delay separates it from the actuator command. Wheel encoders and gyroscopes deliver data at 100 Hz or more but drift, over-read

during wheel slip and carry bias; camera and LiDAR localisation produce absolute pose fixes at 10 Hz or less, with tens of milliseconds of latency and occasional gross errors [1–3]. Raw readings therefore cannot be fed to the controller, and the standard remedy is a recursive estimator [4–6].

Two aspects deserve a controlled, quantitative treatment. First, delayed measurements: visual–inertial and LiDAR–inertial estimators handle timing inside their own back ends [7–10], but in a simple filter built around wheel odometry a pose fix that arrives 60 ms late is easily applied to the current state as if it were fresh. Exact treatments exist [11, 12] and filter packages offer a history-based variant [13], but the effect of delay handling, innovation gating and calibration offsets on closed-loop tracking is not isolated in the works cited here. Second, adaptation: adaptive control addresses parametric change online [14–16], but stability results are usually derived for exact state feedback in continuous time, and it is seldom stated which assumptions survive a filtered state, a fixed sampling period and saturating actuators.

This paper studies both in one simulation: a differential-drive robot follows a figure-eight path while payload and friction change mid-run. All numbers come from the simulation code described in Section 2.4; no hardware experiment was performed. Contributions: (1) a multi-rate extended Kalman filter (EKF) pipeline with timestamp handling, delayed-measurement replay, innovation gating and explicit calibration, with a numerically verified replay; (2) a projection-based adaptive term added to a PID, an analysis of which boundedness properties follow from stated assumptions and which do not, and a safety filter with deadline logic; (3) a pre-specified study with 40 paired seeds that quantifies what each component contributes, including where it does not help, and a wall-clock timing of the code. Wheeled-robot modelling follows [17]; model predictive control [18], dynamic-window avoidance [19] and the wider planning-and-control landscape [20] are outside the scope, so that estimation and adaptation can be isolated.

2. Materials and Methods

2.1 System model, sensors and calibration

The pose is (p_x, p_y, θ) with forward speed v and yaw rate ω; the kinematics are the unicycle model [21, 22]. For $j \in \{v, \omega\}$ the simulated plant has first-order dynamics with viscous friction, actuator delay $\tau_u = 20$ ms, command $|u_j| \le 1$ and white acceleration noise w_j (0.03 m $s^{-3/2}$, 0.06 rad $s^{-3/2}$):

$$\dot{p}_x = v\cos\theta, \quad \dot{p}_y = v\sin\theta, \quad \dot{\theta} = \omega, \quad \dot{y}_j = -a_j y_j + b_j u_j(t - \tau_u) + w_j, \quad y_j \in \{v, \omega\} \tag{1}$$

Nominal values are $a_v = 2\ s^{-1}$, $b_v = 3\ m\ s^{-2}$, $a_\omega = 4\ s^{-1}$, $b_\omega = 12\ rad\ s^{-2}$. At $t = 20$ s a payload step changes $b_j \to 0.6\, b_j$ and $a_j \to 1.5\, a_j$ in both channels, unknown to the controller design. Sensors are proprioceptive (encoders, gyroscope), exteroceptive (camera, LiDAR) or navigational (GNSS, markers; not simulated); each obeys $\mathbf{z}_i = \mathbf{h}_i(\mathbf{x}) + \mathbf{v}_i$ with $\mathbf{v}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_i)$ (Table 1). The estimator state is $\mathbf{x} = [p_x, p_y, \theta, v, \omega, b_g]^\top$ with b_g the gyroscope bias.

Table 1. Simulated sensors and error models.

Sensor	Observes	Rate	Error model in the simulation
Wheel encoders	v, ω	100 Hz	noise 0.01 m/s, 0.03 rad/s; wheel-scale error N(0, 1 %) per run; slip events (0.15 s^{-1}, 0.4 s): speed over-read 15–40 %, yaw offset ±0.4 rad/s
Gyroscope (IMU)	$\omega + b_g$	100 Hz	noise 0.005 rad/s; initial bias 0.005–0.010 rad/s; random walk $2 \cdot 10^{-4}$ rad $s^{-3/2}$
Camera / LiDAR pose fix	(p_x, p_y, θ) of sensor frame	10 Hz	noise 0.05 m, 0.03 rad; latency 60 ± 10 ms; 2 % outliers (4 sigma per component, random sign); extrinsic error 1 cm, 0.3°

Each sensor frame relates to the base frame by a planar rigid transform $\mathbf{T} = (l_x, l_y, \psi)$; the extrinsics $\mathbf{T}_{\text{lidar}}^{\text{cam}}$ and $\mathbf{T}_{\text{imu}}^{\text{base}}$ are calibrated offline and compose as $\mathbf{T}_{\text{lidar}}^{\text{base}} = \mathbf{T}_{\text{cam}}^{\text{base}} \mathbf{T}_{\text{lidar}}^{\text{cam}}$. A pose fix is then a function of the base pose through the mounting offset:

$$\mathbf{h}_{\text{cam}}(\mathbf{x}) = \begin{bmatrix} p_x + l_x \cos\theta - l_y \sin\theta \\ p_y + l_x \sin\theta + l_y \cos\theta \\ \theta + \psi \end{bmatrix} \tag{2}$$

The true offset is $(0.15\ \text{m}, 0.05\ \text{m}, 0.10\ \text{rad})$; the estimator uses a value off by a fixed random error (1 cm, 0.3°, drawn per seed) that is deliberately not modelled in the filter. Fig. 1 shows the pipeline. Every packet carries its acquisition timestamp, and the estimator buffers state, covariance, packet and applied command for the last 32 steps.

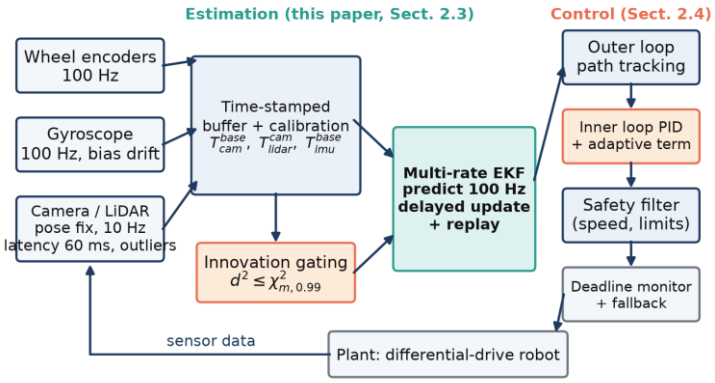

Fig. 1. Pipeline architecture. Encoder, gyroscope and delayed camera or LiDAR pose fixes pass through a time-stamped, calibrated buffer and an innovation gate into a multi-rate EKF with replay; the estimate drives the outer tracking loop, the inner PID with adaptive term, the safety filter and the deadline monitor.

2.2 Multi-rate EKF with delayed measurements and gating

Prediction integrates the unicycle kinematics and the nominal velocity dynamics driven by the known (delayed) applied command, with $\Delta t = 10$ ms:

$$\hat{\mathbf{x}}_{k|k-1} = \begin{bmatrix} \hat{p}_x + \hat{v}\cos\hat{\theta}\,\Delta t \\ \hat{p}_y + \hat{v}\sin\hat{\theta}\,\Delta t \\ \hat{\theta} + \hat{\omega}\,\Delta t \\ \hat{v} + (-a_v\hat{v} + b_v u_v)\Delta t \\ \hat{\omega} + (-a_\omega\hat{\omega} + b_\omega u_\omega)\Delta t \\ \hat{b}_g \end{bmatrix}_{k-1|k-1} \tag{3}$$

so after the payload step the prediction is wrong by design and the process noise must absorb the mismatch. With the Jacobian $\mathbf{F}_k$ (identity plus $-\hat{v}\sin\hat{\theta}\,\Delta t$ and $\hat{v}\cos\hat{\theta}\,\Delta t$ in the heading column of the position rows, $\cos\hat{\theta}\,\Delta t$ and $\sin\hat{\theta}\,\Delta t$ in the speed column, Δt in the yaw-rate column of the heading row, and $1 - a_v\Delta t$, $1 - a_\omega\Delta t$ on the velocity diagonal):

$$\mathbf{P}_{k|k-1} = \mathbf{F}_k\,\mathbf{P}_{k-1|k-1}\,\mathbf{F}_k^{\top} + \mathbf{Q}, \qquad \mathbf{Q} = \mathrm{diag}(q_p, q_p, q_\theta, q_v\Delta t, q_\omega\Delta t, q_b\Delta t) \tag{4}$$

with $q_p = 10^{-7}$, $q_\theta = 10^{-8}$, $q_v = 0.64$, $q_\omega = 2.25$, $q_b = 4 \cdot 10^{-8}$; the velocity entries are far above the physical noise because they cover the model mismatch (selected on separate tuning seeds). The encoder-gyroscope packet $[v_e, \omega_e, \omega_g]^{\top}$ updates the filter every cycle through a linear $\mathbf{H}_p$. A generic update with $\mathbf{z}$, $\mathbf{H}$, $\mathbf{R}$ uses the Joseph form [5, 23]:

$$\begin{aligned} \mathbf{S} &= \mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^{\top} + \mathbf{R}, \qquad \mathbf{K} = \mathbf{P}_{k|k-1}\mathbf{H}^{\top}\mathbf{S}^{-1}, \\ \hat{\mathbf{x}}_{k|k} &= \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}\mathbf{y}, \\ \mathbf{P}_{k|k} &= (\mathbf{I} - \mathbf{K}\mathbf{H})\mathbf{P}_{k|k-1}(\mathbf{I} - \mathbf{K}\mathbf{H})^{\top} + \mathbf{K}\mathbf{R}\mathbf{K}^{\top} \end{aligned} \tag{5}$$

where $\mathbf{y} = \mathbf{z} - \mathbf{h}(\hat{\mathbf{x}}_{k|k-1})$ (heading innovations are wrapped). For a pose fix $\mathbf{h} = \mathbf{h}_{cam}$ and $\mathbf{H}_{cam}$ is its Jacobian. The unscented transform [24] would avoid linearisation, but the nonlinearity is mild. **Gating.** A packet is accepted only if the squared Mahalanobis distance passes a chi-square test with $m = 3$ degrees of freedom (threshold 11.34, 99 %), which by construction rejects 1 % of valid packets and cannot detect small outliers:

$$d^2 = \mathbf{y}^{\top}\mathbf{S}^{-1}\mathbf{y} \le \chi^2_{m,\,0.99} \tag{6}$$

Replay. A fix acquired at step m arrives at step $k > m$ (5 to 7 steps). Applying it to $\hat{\mathbf{x}}_{k|k-1}$ as if fresh misplaces the robot by $\hat{v}\,\tau$. The replay variant restores the stored posterior $(\hat{\mathbf{x}}_{m|m}, \mathbf{P}_{m|m})$, applies the gated update at t_m and re-runs prediction $\mathcal{P}$ and the stored gated updates $\mathcal{U}$:

$$(\hat{\mathbf{x}}_{j|j}, \mathbf{P}_{j|j}) = \mathcal{U}\big(\mathcal{P}\big(\hat{\mathbf{x}}_{j-1|j-1}, \mathbf{P}_{j-1|j-1}, \mathbf{u}_j\big), \mathbf{z}_{p,j}\big), \quad j = m+1, \ldots, k \tag{7}$$

This reproduces the estimate that would have been obtained had the fix been available at t_m (verified in Section 3.1) at the cost of $k - m$ extra filter steps per fix; recursions without a stored history [11, 12] are cheaper in memory. Consistency is monitored with the normalised estimation error squared (NEES) of the pose and the normalised innovation squared (NIS), both chi-square with 3 degrees of freedom for a consistent filter [5, 25]:

$$\mathrm{NEES}_k = \tilde{\mathbf{e}}_k^{\top}\mathbf{P}^{-1}_{\mathrm{pose},k|k}\,\tilde{\mathbf{e}}_k, \qquad \tilde{\mathbf{e}}_k = [p_x - \hat{p}_x,\ p_y - \hat{p}_y,\ \theta - \hat{\theta}]^{\top} \tag{8}$$

2.3 Adaptive control, safety filter and deadline logic

The outer layer is the kinematic tracking law of Kanayama et al. [26] on the estimated pose; with robot-frame errors e_x, e_y, e_θ and fixed gains $k_x = 1.2$, $k_y = 2.0$, $k_\theta = 2.5$ it yields desired speeds

$$v_d = v_{ref}\cos e_\theta + k_x e_x, \qquad \omega_d = \omega_{ref} + v_{ref}\left(k_y e_y + k_\theta \sin e_\theta\right) \tag{9}$$

The inner layer tracks $d_j \in \{v_d, \omega_d\}$ with error $e_j = d_j - \hat{y}_j$ and realises $u = u_0 + \Delta u_{ad}$ as a PID law with static feedforward (nominal parameters) plus an adaptive term:

$$u_j = \frac{a_j^{nom}}{b_j^{nom}} d_j + K_{p,j}\, e_j + K_{i,j} \int e_j\, dt + K_{d,j}\, \dot{e}_j + \boldsymbol{\theta}_j^{\top} \boldsymbol{\phi}_j, \qquad \boldsymbol{\phi}_j = [\, d_j, \dot{d}_j \,]^{\top} \tag{10}$$

For first-order dynamics there is an ideal $\boldsymbol{\theta}_j^{id} = [a_j/b_j, 1/b_j]^{\top}$ giving the exact inverse-dynamics command. The parameters follow a normalised gradient law with projection onto the box $\Omega = \{|\theta_i| \le 2\}$ (componentwise clipping) and a dead zone $e_0 = 0.005$:

$$\boldsymbol{\theta}_j(k+1) = \mathrm{Proj}_{\Omega}\left\{ \boldsymbol{\theta}_j(k) + \frac{\Gamma_j\, \boldsymbol{\phi}_j(k)\, e_j(k)\, \Delta t}{1 + \|\boldsymbol{\phi}_j(k)\|^2} \right\} \text{ if } |e_j(k)| > e_0 \tag{11}$$

with $\Gamma_v = 30$, $\Gamma_\omega = 15$; the sign of b_j is assumed known. The integrator uses conditional anti-windup. **Boundedness argument** [14–16, 27]. For one channel in continuous time with unknown constant a, $b > 0$, proportional feedback K_p and the adaptive term (integral and derivative dropped), with $\tilde{\boldsymbol{\theta}} = \boldsymbol{\theta} - \boldsymbol{\theta}^{id}$ and a disturbance w that includes the estimation error, $\dot{e} = -(a + bK_p)e - b\tilde{\boldsymbol{\theta}}^{\top}\boldsymbol{\phi} - w$. With

$$V = \frac{1}{2}e^2 + \frac{b}{2\Gamma}\tilde{\boldsymbol{\theta}}^{\top}\tilde{\boldsymbol{\theta}}, \qquad \dot{V} \le -(a + bK_p)\, e^2 + |e|\,\overline{w} \quad (|w| \le \overline{w}) \tag{12}$$

(the cross terms cancel, and projection onto a convex Ω that contains $\boldsymbol{\theta}^{id}$ can only decrease $\tilde{\boldsymbol{\theta}}^{\top}\dot{\boldsymbol{\theta}}$), V decreases whenever $|e| > \overline{w}/(a + bK_p)$; $\tilde{\boldsymbol{\theta}}$ is bounded by projection, so e is uniformly ultimately bounded by $\sqrt{(\overline{w}/(a + bK_p))^2 + (b/\Gamma)D_\Omega^2}$ (D_Ω: diameter of Ω) if $\boldsymbol{\phi}$ is bounded. Parameter convergence is not claimed (it needs persistent excitation).

> **Scope.** This is proven only for the scalar, continuous-time, full-measurement channel. It is not proven for the implemented system, which differs by (i) discrete updates with normalisation and dead zone, (ii) PID integral and derivative terms, (iii) saturation and a 20 ms delay, (iv) coupling through the outer loop, whose input-to-state stability is not shown, (v) feedback from an EKF whose error bound $\overline{w}$ is assumed, not derived, and (vi) jumping plant parameters. The adaptive controller is demonstrated by simulation and a numerical check, not proven stable.

Safety filter. The speed is limited so the robot can stop, at $a_b = 0.8\ \mathrm{m\,s^{-2}}$, before a margin μ around a mapped circular obstacle (radius 0.25 m; robot radius 0.20 m), if the obstacle lies within ±75° of the estimated heading; c is the estimated clearance:

$$v \le \sqrt{2\, a_b \max(c - \mu,\, 0)}, \qquad \mu = \mu_0 + k_\sigma\, \sigma_p, \quad \sigma_p = \sqrt{\lambda_{\max}(\mathbf{P}_{xy})} \tag{13}$$

The fixed variant uses $\mu_0 = 0.10$ m, $k_\sigma = 0$; the covariance-aware variant $k_\sigma = 3$, so the robot slows when uncertainty grows. This is a speed limiter, not a control-barrier-function program [28], and gives no guarantee when the estimate is wrong. **Deadline logic.** If the cycle exceeds T_s the previous command is held; after five consecutive misses the command is zero. Algorithm 1 summarises the cycle.

Algorithm 1. Real-time sensor-fusion and adaptive-control cycle

1. Read packets with acquisition stamps; apply calibration transforms; predict the EKF with the applied command; gate and apply the encoder-gyroscope update; store the posterior.
2. If a pose fix with stamp m arrived: restore the posterior at m, apply the gated fix, replay steps $m + 1, \ldots, k$ with the stored inputs.
3. Outer loop: (v_d, ω_d) from the reference and $\hat{\mathbf{x}}_k$; safety filter: $v_d \leftarrow \min(v_d, v_{\max}(c))$ using σ_p.
4. Inner loop: PID with feedforward plus adaptive term, projection and dead zone; clip and anti-windup.
5. If the elapsed time exceeds T_s, hold the previous command (zero after 5 consecutive misses); otherwise transmit. Log latency, e, **P**, saturation, deadline status.

2.4 Simulation design, metrics and statistics

The reference is a figure-eight (Gerono lemniscate, $A = 2.5$ m, length 15.24 m) at 0.5 m/s (lap 30.5 s, largest yaw rate 0.96 rad/s); runs last 60 s at 10 ms with the payload step at 20 s. Estimators: **E0** dead reckoning; **E1** EKF with encoders and gyroscope only (no fixes, no gating); **E2** full pipeline; **A1** E2 without delay handling; **A2** E2 without gating. Controllers on the estimated state: **C1** fixed PID; **C2** C1 plus the adaptive term; **C3** fixed model-based inverse dynamics $u_j = (\dot{d}_j + a_j^{nom} d_j + \lambda_j e_j)/b_j^{nom}$ without integral action. MPC was not simulated. All code is Python (NumPy). Every random input is generated per seed before the run, so conditions with the same seed share noise and differences are paired. Tuning (EKF noise from five candidate pairs, seeds 1000–1005; PID gains, adaptation rates and C3 gain from four, four and three candidates, seeds 1000–1009; rule: lowest 20–30 s tracking error among settings with actuator saturation below 7 %) used seeds disjoint from the test seeds 0–39 ($n = 40$), and nothing was changed afterwards. Scenarios: the main figure-eight; an obstacle on the path at the point reached at 12 s (C2; E0, E1, E2; no filter, fixed margin, covariance margin); synthetic controller deadline misses (0–20 % of cycles) and single outages (0.03–1 s at 25 s); $T_s = 20$ ms; and a run without payload change.

Table 2 defines the metrics. Tracking metrics use the true pose; the conventional e_{RMS} on the estimated state is also computed but not compared. Pre-specified primary comparisons: P1 E1 vs E2, P2 A1 vs E2, P3 A2 vs E2 (position RMSE); P4 C1 vs C2 and P5 C3 vs C2 (post-step e_{RMS}, 20–60 s, under E2); P6 C2 vs C1 before the step (harm check, 5–20 s); P7 covariance-aware safety filter vs none (E2, C2). Others are secondary. Each paired difference is reported as the mean with a Student-t interval, with Bonferroni-adjusted intervals (level $1 - 0.05/7$) for P1–P7 and a bootstrap check:

$$\overline{\Delta} \pm t_{0.975,\,n-1}\,\frac{s_\Delta}{\sqrt{n}}, \qquad n = 40 \tag{14}$$

Table 2. Metrics.

Metric	Definition	Unit
Tracking e_{RMS}	$\sqrt{\frac{1}{N}\sum_i \|\mathbf{p}_{ref}(t_i) - \mathbf{p}(t_i)\|^2}$, true position; 5–20 s (pre), 20–60 s (post), 20–30 s (transient)	m
Estimator RMSE	RMS of $\|\mathbf{p} - \hat{\mathbf{p}}\|$ and wrapped heading error, 2–60 s	m, deg
NEES, NIS; gate rates	mean NEES (every 50 ms) and NIS over valid fixes (expected 3); share of outliers and of valid fixes rejected	-, %
Speed-loop deviation	$e_v = v - v_d$ smoothed over 0.2 s: peak, positive overshoot; recovery time (time after 20 s from which $\|e_v\| \le 0.05$ m/s for 1 s; 40 s if never); IAE over 20–30 s	m/s, s, m
Effort, saturation	RMS of u_v, u_ω; share of cycles at the command limit	-, %
Clearance	minimum true distance between robot and obstacle surfaces; contact if negative	m
Compute time, deadline	wall-clock time per cycle of estimator, controller and safety filter (plant excluded): mean, p95, p99, max; share of cycles within T_s (sums of two cycles for 20 ms)	ms, %

Verification (Section 3.1): analytic Jacobians against central differences at 200 random states; Joseph form against the textbook update; the replay filter against an oracle filter applying each fix at its acquisition time (with and without gating); the scalar adaptive argument by simulation; NEES and NIS in the main study. **Use of AI:** large language models (Claude, Anthropic) were used under the authors' direction for language editing, structuring and for writing and running the simulation code (see Declarations).

3. Results

All results are **simulation** results (seeds 0–39); none is a measurement on a physical robot.

3.1 Verification

Jacobians agreed with finite differences to $2.0 \cdot 10^{-10}$ ($\mathbf{F}_k$) and $3.2 \cdot 10^{-10}$ ($\mathbf{H}_{cam}$); the Joseph form matched the textbook update to $1.3 \cdot 10^{-19}$ (relative). The replay filter reproduced the oracle with a maximum absolute difference of 0 (bitwise identical) with and without gating, whereas applying the fix to the current state differed by up to 0.021 in state. In the scalar adaptive check (a: 2→3, b: 3→1.8 at 20 s, $K_p = 3$, $\Gamma = 30$) the parameters stayed inside the box (largest magnitude 1.69 against 2) and V never increased without disturbance; with a 0.05 sinusoidal disturbance the largest error was 0.0073 against the decrease threshold $\overline{w}/(a + bK_p) = 0.0060$, slightly above it, as expected, because the guaranteed bound includes the parameter-error term. These checks verify the implementation, not the system against a robot.

3.2 State estimation

Table 3 and Fig. 2 compare the estimators under C2. The full pipeline E2 reached a position RMSE of 0.024 m (SD 0.007) and a heading RMSE of 0.57°, against 0.372 m and 10.50° for E1 and 0.488 m and 10.52° for dead reckoning E0; without fixes the error accumulates (median 0.1–0.8 m in Fig. 2b). E0 and E1 have the same heading error: the gyroscope alone, without gating or absolute heading, did not help, most likely because slip corrupts the encoder yaw rate against which its bias is estimated.

Table 3. Estimator comparison under C2 (simulation, 40 seeds; mean ± SD). Expected NEES and NIS: 3. E0 dead reckoning; E1 EKF without fixes; E2 full pipeline; A1 without delay handling; A2 without gating. Counts pooled over seeds.

Est.	Position RMSE (m)	Heading RMSE (°)	NEES [95 % CI]	NIS valid fixes [95 % CI]	Outliers rejected	Valid rejected
E0	0.488 ± 0.171	10.52 ± 4.51	-	-	-	-
E1	0.372 ± 0.143	10.50 ± 4.48	110.5 [85.6, 135.4]	-	-	-
E2	0.024 ± 0.007	0.57 ± 0.12	10.2 [8.2, 12.3]	3.15 [3.08, 3.22]	499 / 499	1.3 %
A1	0.035 ± 0.006	1.35 ± 0.15	32.0 [28.9, 35.1]	3.42 [3.34, 3.50]	499 / 499	2.1 %
A2	0.038 ± 0.010	1.16 ± 0.30	48.1 [42.3, 54.0]	3.98 [3.84, 4.11]	0 / 499	0.0 %

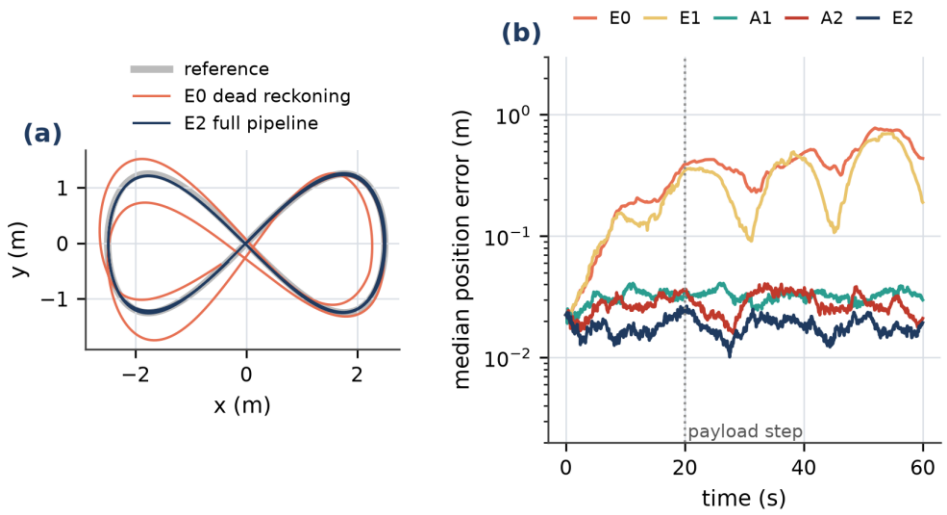

Fig. 2. (a) Reference path and true trajectories of seed 0 under dead reckoning (E0) and the full pipeline (E2), both with C2. (b) Median over 40 seeds of the position estimation error for the five estimators; the dotted line is the payload step.

Paired differences (P1–P3; Fig. 3b): adding fixes (E1 minus E2) lowered position RMSE by 0.348 m (95 % CI 0.303 to 0.393; E1 error 15 times E2). Replacing replay by naive application (A1 minus E2) raised it by 0.0109 m (95 % CI 0.0084 to 0.0134; 45 %) and heading RMSE by 0.77° (95 % CI 0.71 to 0.84). Removing gating (A2 minus E2) raised it by 0.0137 m (95 % CI 0.0117 to 0.0158; 57 %). All three intervals exclude zero after Bonferroni adjustment, and bootstrap intervals agree (P2: 0.0085 to 0.0132 m). The gate rejected 499 of 499 outliers in

E2 and 1.3 % of valid fixes (1 % expected); A1 rejected 2.1 %. Consistency is only partly met: NIS of valid fixes was 3.15 (95 % CI 3.08 to 3.22) in E2, near 3, but NEES was 10.2 (95 % CI 8.2 to 12.3) (Fig. 3a): **the filter is overconfident**, because unmodelled wheel-scale error, extrinsic error and slip passing the gate are small against the fix noise but dominate the covariance once fixes have shrunk it. NEES was 32.0 for A1, 48.1 for A2 and 110 for E1.

3.3 Tracking after the payload step

Table 4 compares the controllers under E2 (MPC was not simulated). Before the step all were indistinguishable (about 0.027 m). After it, C3, which has no integral action, rose to 0.103 m, whereas C1 and C2 stayed at 0.038 m. P4 (C1 minus C2, post-step e_{RMS}) was 0.0002 m (95 % CI −0.0004 to 0.0008): **the adaptive term did not change the overall tracking error**. P6 (C2 minus C1 before the step) was −0.0002 m (95 % CI −0.0004 to −0.0001), a negligible 0.2 mm, so no harm. Against C3 (P5) C2 was better by 0.065 m (95 % CI 0.062 to 0.067), mostly reflecting C3's missing integral action. In secondary speed-loop measures (Fig. 4a) the term helped: 20–30 s tracking error fell by 0.0025 m (95 % CI 0.0016 to 0.0035), speed-loop IAE by 0.032 m (0.028 to 0.037) and recovery time by 0.33 s (0.25 to 0.41). The costs: saturation rose by 3.9 percentage points (95 % CI 3.4 to 4.4), and peak speed-loop deviation (0.237 vs 0.208 m/s) and overshoot (6.1 % vs 4.1 %) were larger. Secondary comparisons are unadjusted and descriptive.

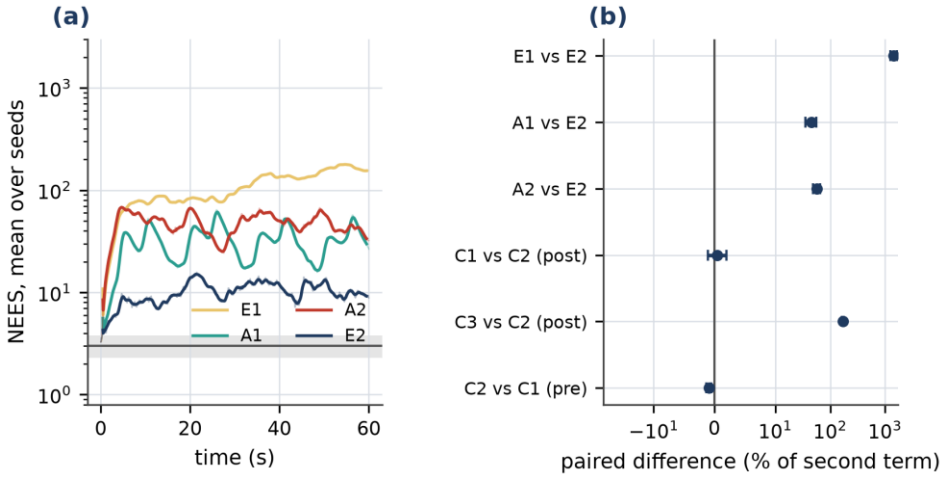

Fig. 3. (a) NEES of the pose, mean over 40 seeds (thin: raw; thick: 1 s average); black line: expected value 3; grey band: its 95 % interval. (b) Paired differences for P1–P6 as a percentage of the second term with 95 % confidence intervals (symmetric-logarithmic axis); positive means the first term is larger.

Table 4. Controller comparison under E2 (simulation, 40 seeds; mean ± SD). Recovery is 40 s if the speed deviation never stays within 0.05 m/s for 1 s; clearance is from the obstacle scenario with the covariance-aware filter.

Controller (with E2)	e_{RMS} 5–20 s (m)	e_{RMS} 20–60 s (m)	Recovery (s)	Speed IAE (m)	Sat. (%)	Min. clear. (m)
C1 fixed PID	0.027 ± 0.009	0.038 ± 0.012	1.02 ± 0.69	0.224 ± 0.101	2.4	0.081
C2 PID + adaptive	0.027 ± 0.009	0.038 ± 0.012	0.68 ± 0.71	0.192 ± 0.102	6.3	0.094
C3 fixed model-based	0.026 ± 0.009	0.103 ± 0.009	40.00 ± 0.00	1.120 ± 0.095	3.8	0.092

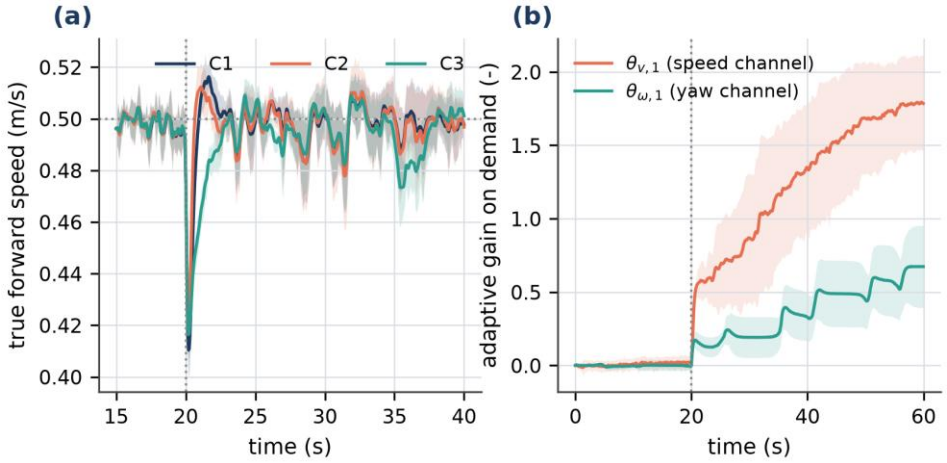

Fig. 4. Payload step at 20 s (simulation, E2). (a) True forward speed, mean over 40 seeds with 95 % band of the mean, for C1, C2 and C3; dotted: 0.5 m/s demand. (b) Adaptive gains on the desired value in the speed and yaw channels (mean ± SD) for C2.

Without a payload change, errors were 0.0285 m (C1), 0.0282 m (C2) and 0.0278 m (C3). The adaptive parameters did not converge: $\theta_{v,1}$ jumped to about 0.6 within a second of the step and then drifted to 1.78 (SD 0.32) at 60 s, above the value 1.0 that would make the feedforward exact and near the bound 2 (Fig. 4b). The PID integrator already removes the steady-state error, the speed regressor lacks persistent excitation and the dead zone is below the noise level; this is the known drift of gradient laws without leakage, and projection keeps it bounded. Under E0 or E1 the C1–C2 difference was 0.0001 and −0.0004 m, dominated by estimator drift.

3.4 Obstacle scenario

Without the filter the robot hit the obstacle in all 40 runs for every estimator, so the filter decides the outcome (Fig. 5a). With E2 no run made contact (0 of 40) under either margin; the covariance margin kept more clearance (mean 0.094 m vs 0.058 m; paired difference 0.036 m, 95 % CI 0.033 to 0.039). The filter is only as good as the estimate: with the fixed margin, E0 and E1 made contact in 19 and 21 of 40 runs, and the covariance margin reduced E1 to 11 of 40 without removing it (clearance difference 0.099 m, 95 % CI 0.095 to 0.103). P7 (covariance filter minus none, E2) was 0.534 m (95 % CI 0.531 to 0.538; positive means more clearance).

3.5 Computation time, deadlines and sampling

The per-cycle wall-clock time of the estimator, controller and safety filter was measured in Python 3.14.4 (NumPy 2.4.4) on 11th Gen Intel(R) Core(TM) i3-1115G4 @ 3.00GHz

(Windows-11-10.0.26200-SP0), 30 further seeds (100–129), 180,000 cycles per configuration, on an otherwise idle machine with configurations interleaved per seed (an earlier measurement taken while other jobs ran had a heavier tail and was discarded). This is a real measurement of this code on this PC under a non-real-time operating system, not a worst-case execution time; embedded targets will differ. For E2 + C2 (Table 5, Fig. 5b) the mean was 0.361 ms, the 99th percentile 1.55 ms and the maximum 41.7 ms; cycles with a replayed fix cost more (mean 1.07 ms vs 0.28 ms). Compliance was 99.83 % for 10 ms (297 misses) and 99.91 % for 20 ms (sums of two cycles); the misses are isolated stalls of the operating system and interpreter, not the algorithm. Lengthening the control period to 20 ms raised C2's tracking error from 0.0357 to 0.0390 m (paired difference 0.0033, 95 % CI 0.0026 to 0.0041). Synthetic misses of up to 20 % of cycles changed it by only 0.0001 m (95 % CI −0.0000 to 0.0003), since holding a command for a few 10 ms cycles is harmless for a speed loop with a time constant near 0.3 s; this shows that the fallback logic runs without instability, not that it protects against long outages. Single outages (0.03–1 s) changed post-step errors only for 1 s (to 0.114 and 0.112 m for C1 and C2, a stopped robot in front of a moving reference).

Table 5. Measured per-cycle compute time in ms (Python, test PC, 30 seeds × 6000 cycles; plant simulation excluded) and deadline compliance for 10 ms and 20 ms (sums of two cycles).

Configuration	Mean	Median	p95	p99	Max	10 ms (%)	20 ms (%)
E2 + C1	0.311	0.201	0.886	1.419	37.8	99.87	99.90
E2 + C2	0.361	0.245	0.942	1.547	41.7	99.83	99.91
E2 + C3	0.286	0.180	0.854	1.327	62.8	99.87	99.90
E1 + C2	0.279	0.237	0.384	0.822	60.5	99.88	99.93

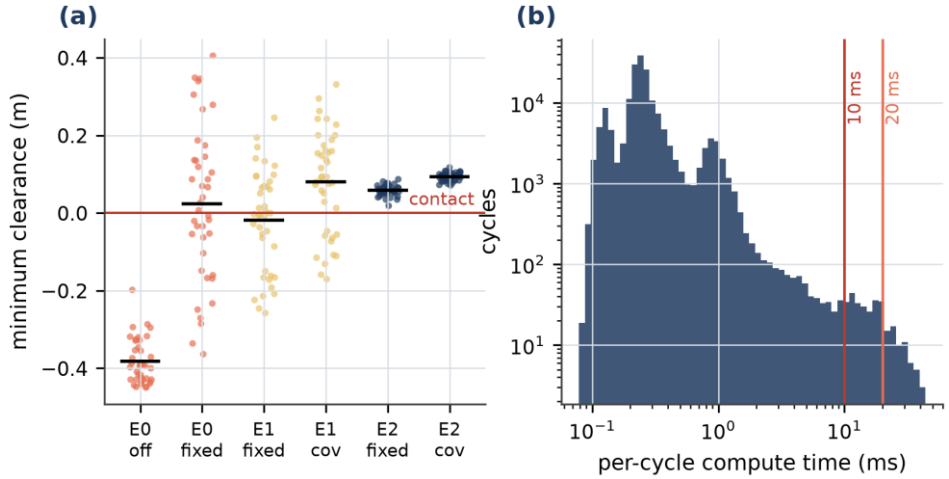

Fig. 5. (a) Minimum true clearance in 40 runs for six estimator and filter combinations (points: seeds; bars: means; red line: contact). (b) Distribution of per-cycle compute time of E2 + C2, log-log axes, with the 10 ms and 20 ms periods.

4. Discussion

In this simulation estimator design mattered far more than adaptation. The full pipeline cut the position error by 94 % relative to the encoder-gyroscope filter, and replay (45 %) and gating

(57 %) each gave a clear gain, consistent with the known benefit of exact treatment of delayed data [11, 12]. The delay here (60 ms at 0.5 m/s, 3 cm) is small against the fix noise (5 cm); the benefit would grow with speed and latency, and 2–3 sigma or correlated outliers, where gating is expected to fail, were not tested. The adaptive term did not improve the primary tracking metric: a PID with an integrator already compensates a gain and friction step, so adaptation acted mostly in the first seconds, shortening recovery at the price of more saturation and a larger transient. This is not evidence that adaptation is useless; the plant is first-order, the change is a single bounded step and the baseline is well tuned on the same model structure. Practically, replay is cheap (mean cycle 0.28 to 1.07 ms when a fix arrived); a safety filter needs an estimate with bounded error, which a margin cannot replace; and the 10 ms deadline was met in 99.83 % of cycles on a laptop-class processor, but embedded targets need worst-case execution-time analysis [29–31] and a real-time middleware configuration [32]. EKF overconfidence is a recognised property of EKF localisation and SLAM [25, 33, 34]; here its dominant cause is the deliberate omission of calibration and wheel-scale errors, so the result concerns unmodelled systematic errors rather than linearisation.

Limitations. (1) Simulation only: real slip, sensor failure and unmodelled dynamics would change the numbers. (2) The filter and controller share the plant's model structure (first-order velocity dynamics), which favours both; second-order actuators or Coulomb friction would test them harder. (3) One path, step size and obstacle geometry; seeds vary noise, slip, bias and calibration, not plant structure. (4) Settings were tuned on separate seeds but on the same simulator, so the closeness of C1 and C2 may reflect good tuning of C1. (5) The boundedness proof covers only the scalar continuous-time channel. (6) Only 4-sigma outliers were simulated, and the map and obstacle were assumed known. (7) The timing is for interpreted Python on a desktop-class CPU and gives no worst-case bound. (8) The 40 seeds share one plant and scenario, so the intervals describe noise-realisation variation only. Future work: a physical platform with logged timestamps, structurally different plants, smaller and correlated outliers, a discrete-time stability analysis of the cascade with the estimator in the loop, and an MPC baseline [18].

5. Conclusions

A multi-rate EKF with timestamp replay of delayed fixes, chi-square gating and explicit calibration kept the simulated position error at 0.024 m, against 0.372 m without fixes and 0.488 m for dead reckoning; omitting delay handling or gating degraded it by 45 % and 57 %. A projection-based adaptive term shortened speed recovery after a payload and friction step but did not improve overall tracking and increased saturation. The safety filter prevented contact only under the full pipeline. The boundedness argument covers a scalar continuous-time channel; the implemented system was verified numerically and in simulation, not by proof or on hardware, and real-robot validation remains the necessary next step.

Declarations

Author contributions (CRediT). Both authors contributed equally to conceptualization, methodology, software, validation, formal analysis, visualization, and writing (original draft, review and editing). **Funding.** No external funding. **Competing interests.** None declared. **Data**

and code availability. Simulation scripts and generated data are available from the corresponding author on request. **Use of AI.** Large language models (Claude, Anthropic; models Claude Fable 5.1, Claude Sonnet 5.5 and Claude Opus 5.5, accessed through Claude Code) were used under the authors' direction for language editing, structuring of the manuscript, drafting of figures, and for writing and running the simulation code reported in Sections 2.4 and 3. The authors reviewed and verified the content and take full responsibility for it.

References

1. U. Mayinova and J. M. Abdiev, “Enhancing energy efficiency in industrial enterprises through advanced automation and control systems,” JizPI Xabarnomasi, vol. 1, no. 1, pp. 989–992, 2024.
2. Siegwart, R., Nourbakhsh, I. R., Scaramuzza, D.: Introduction to Autonomous Mobile Robots, 2nd ed. MIT Press, Cambridge, MA (2011).
3. Abdiev, J. (2026a). Beyond faradaic efficiency: A closure audit of charge, carbon utilization, and energy reporting in electrochemical CO2 reduction. Journal of Electroanalytical Chemistry, 1022, Article 120675. https://doi.org/10.1016/j.jelechem.2026.120675.
4. Abdiev, J. (2026b). Data for: Reassessing isospin-resolved decay of the ^{31}P giant dipole resonance: Covariance-aware channel competition and cascade closure from archival γ-X-γ′ data [Data set]. Mendeley Data. https://doi.org/10.17632/538h8zxwxn.1.
5. Abdiev, J. M. (2022). Physical terms problems and their solutions [Working paper]. SSRN. https://doi.org/10.2139/ssrn.4024965.
6. Abdiev, J., Abdieva, N. M., & Khasanova, D. Y. (2022). Physical terms, problems and their solutions [Working paper]. SSRN. https://doi.org/10.2139/ssrn.4066756.
7. Abdiev, J., & Malikov, R. (2022). The role of teacher and student cooperation in the modern educational process of Google Docs. Web of Scientist: International Scientific Research Journal, 3(1), 579–589.
8. Abdiev, J., & Safarov, O. (2022). Basalt fiber-basic (primary) concepts. PhilPapers. https://philpapers.org/rec/ABDBF.
9. Abdiev, J., Safarov, O., & Julanov, H. (2022). Study of the properties of polymer composites: Reinforcement based on glass and basalt fibers. Eurasian Scientific Herald, 7, 77–88.
10. Abdiev, J. M., Leyderman, A. Yu., Saidov, A. S., & Usmonov, Sh. N. (2021, April 21). Perenos toka v poluprovodnikovykh p-n-strukturakh s bazovoy n-oblasti iz varizonnogo poluprovodnika [Current transport in semiconductor p-n structures with a base n-region made of a graded-gap semiconductor]. In Yosh olimlar va fizik talabalarning I respublika ilmiy anjumani (YoOFTRIA-I) materiallari [Proceedings of the 1st republican scientific conference of young scientists and physics students (YoOFTRIA-I)] (Vol. 1, pp. 114–117). Tashkent. (in Russian).
11. Abdiev, J. M., & Ziyadullaev, Z. R. (2022, April 22). Si taglikka o‘stirilgan (Si2)1-x(GaP)x qattiq qorishma epitaksial qatlamining kimyoviy tarkibi va morfologiyasini o‘rganish [Study of the chemical composition and morphology of the epitaxial layer of the (Si2)1-x(GaP)x solid solution grown on a Si substrate]. In Fizika fanining rivojida iste‘dodli yoshlarning o‘rni (RIAK-XV): Respublika ilmiy-amaliy konferensiyasi ishlari to‘plami

[The role of talented youth in the development of physics (RIAK-XV): Proceedings of the republican scientific-practical conference] (Vol. 2, No. 15, pp. 184–186). (in Uzbek).

12. Abdiyev, J. M., Leyderman, A. Yu., Saidov, A. S., & Usmonov, Sh. N. (2021). Injection diffusion processes in the weak linear graded-band semiconductor p-n-structures. Euroasian Journal of Semiconductors Science and Engineering, 3(1), 20–26. https://doi.org/10.37681/2181-1652-019-x-2021-1-3.
13. Kenjayev, S., & Abdiev, J. (2024). Expanding the role of mathematical modeling in quantum computers: The application of theoretical physics in next-generation computing technologies. Journal of Integrated Education and Research, 3(5), 149–162. https://ojs.rmasav.com/index.php/ojs/article/view/1675.
14. Leyderman, A. Yu., Saidov, A. S., Abdiyev, J. M., Usmonov, Sh. N., & Suyarov, Q. T. (2021, August 15). Influence of weak graded gap on injection diffusion regimes of current transport in semiconductor p-n-structures. In USBIK 2021 Online International Congress on Natural Sciences (Vol. 8, pp. 40–44).
15. Mayinova, U., & Abdiev, J. M. (2024). Enhancing energy efficiency in industrial enterprises through advanced automation and control systems. JizPI Xabarnomasi, 1(1), 989–992.
16. Mustafayev, T., Akhmedzhanov, F., Abdiev, J., & Khasanov, S. (2025a). Correction: An exploration of anisotropic acoustic wave attenuation in quartz crystals. Brazilian Journal of Physics, 55(1), Article 18. https://doi.org/10.1007/s13538-024-01657-x.
17. Mustafayev, T., Akhmedzhanov, F., Abdiev, J., & Khasanov, S. (2025b). An exploration of anisotropic acoustic wave attenuation in quartz crystals. Brazilian Journal of Physics, 55(1), Article 1. https://doi.org/10.1007/s13538-024-01638-0.
18. Payzullaev, A. N., Allaev, B. A., Mirzaev, S. Z., Abdiev, J. M., Urinov, J., & Parkash, A. (2023). The impact of silicon dioxide nanoparticle size on the viscosity and stability of nanofluids: A comprehensive study. ECS Advances, 2(3), Article 031001. https://doi.org/10.1149/2754-2734/ace121.
19. Saidov, A. S., Saparov, D. V., Usmonov, Sh. N., Kutlimratov, A., Abdiev, J. M., Kalanov, M., Razzakov, A. Sh., & Akhmedov, A. M. (2021). Investigation of the crystallographic perfection and photoluminescence spectrum of the epitaxial films of (Si2)1-x(GaP)x ($0 \leq x \leq 1$) solid solution, grown on Si and GaP substrates with the crystallographic orientation (111). Advances in Condensed Matter Physics, 2021, Article 3472487. https://doi.org/10.1155/2021/3472487.
20. Saidov, A. S., Usmonov, Sh. N., Karshiev, A. B., & Abdiev, J. M. (2022). Influence of the varizonal Si1-xGex solid solution composition on the thermovoltaic effect in n-Si-p-(Si)1-xGex structure. IOP Conference Series: Earth and Environmental Science, 1112(1), Article 012040. https://doi.org/10.1088/1755-1315/1112/1/012040.
21. Urinov, Kh. O., Jumanov, Kh. A., Khidirov, A. M., Urinov, S. Kh., Abdiyev, J. M., Jumaboyev, T. A., & Eshmirzayev, M. R. (2020). Magnetocaloric effect in polycrystalline cobalt. Journal of Physics: Conference Series, 1515(2), Article 022079. https://doi.org/10.1088/1742-6596/1515/2/022079.
22. Taylanov, N. A., Urinov, S. Kh., & Abdiev, J. M. (2022). A fourth-order Runge–Kutta method for numerical solution of the Kuramoto–Sivashinsky equation.

23. Urinov, J. O., Abdiev, J. M., & Malikov, R. R. (n.d.). FexV3−xO4: Magnetic and electrical properties of spinel.
24. Komilovich, J. L. H., Ugli, A. A. M., Safarov, O., Ugli, A. J. M., & Ugli, Y. A. A. (2022). Understanding the magneto caloric effect in super spin glass cobalt-based nanoparticles.
25. Bailey, T., Nieto, J., Guivant, J., Stevens, M., Nebot, E.: Consistency of the EKF-SLAM algorithm. In: Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS), pp. 3562–3568 (2006). doi:10.1109/IROS.2006.281644.
26. Kanayama, Y., Kimura, Y., Miyazaki, F., Noguchi, T.: A stable tracking control method for an autonomous mobile robot. In: Proc. IEEE Int. Conf. Robotics and Automation (ICRA), pp. 384–389 (1990). doi:10.1109/ROBOT.1990.126006.
27. Slotine, J.-J. E., Li, W.: On the adaptive control of robot manipulators. Int. J. Robot. Res. 6(3), 49–59 (1987). doi:10.1177/027836498700600303.
28. Ames, A. D., Xu, X., Grizzle, J. W., Tabuada, P.: Control barrier function based quadratic programs for safety critical systems. IEEE Trans. Autom. Control 62(8), 3861–3876 (2017). doi:10.1109/TAC.2016.2638961.
29. Liu, C. L., Layland, J. W.: Scheduling algorithms for multiprogramming in a hard-real-time environment. J. ACM 20(1), 46–61 (1973). doi:10.1145/321738.321743.
30. Buttazzo, G. C.: Hard Real-Time Computing Systems: Predictable Scheduling Algorithms and Applications, 3rd ed. Springer, New York (2011). doi:10.1007/978-1-4614-0676-1.
31. Wilhelm, R., Engblom, J., Ermedahl, A., Holsti, N., Thesing, S., Whalley, D., et al.: The worst-case execution-time problem: overview of methods and survey of tools. ACM Trans. Embed. Comput. Syst. 7(3), 1–53 (2008). doi:10.1145/1347375.1347389.
32. Macenski, S., Foote, T., Gerkey, B., Lalancette, C., Woodall, W.: Robot Operating System 2: design, architecture, and uses in the wild. Sci. Robot. 7(66), eabm6074 (2022). doi:10.1126/scirobotics.abm6074.
33. Huang, S., Dissanayake, G.: Convergence and consistency analysis for extended Kalman filter based SLAM. IEEE Trans. Robot. 23(5), 1036–1049 (2007). doi:10.1109/TRO.2007.903811.
34. Li, M., Mourikis, A. I.: High-precision, consistent EKF-based visual-inertial odometry. Int. J. Robot. Res. 32(6), 690–711 (2013). doi:10.1177/0278364913481251.