

Attack Detection in Dynamic Games with Quadratic Measurements

Muyan Jiang and Anil Aswani

Abstract—This paper studies attack detection for discrete-time linear systems with stochastic process noise that produce both a vulnerable (i.e., attackable) linear measurement and a secured (i.e., unattackable) quadratic measurement. The motivating application of this model is a dynamic-game setting where the quadratic measurement is interpreted as a system-level utility or reward, and control inputs into the linear system are interpreted as control policies that, once applied, are known to all game participants and which steer the system towards a game-theoretic equilibrium (e.g., Nash equilibrium). To detect attacks on the linear channel, we develop a novel quadratic-utility-aware observer that leverages the secured quadratic output and enforces measurement consistency via a projection step. We establish three properties for this observer: feasibility of the true state, prox-regularity of the quadratic-constraint set, and a monotone error-reduction guarantee in the noise-free case. To detect adversarial manipulation, we compare linear and quadratic observer trajectories using a wild bootstrap maximum mean discrepancy (MMD) test that provides valid inference under temporal dependence. We validate our framework using numerical experiments of a pursuit–evasion game, where the quadratic observer preserves estimation accuracy under linear-sensor attacks, while the statistical test detects distributional divergence between the observers’ trajectories.

I. INTRODUCTION

Secure state estimation is critical for multi-agent systems in which multiple decision-makers coordinate actions from streamed sensor data [1], [2]. There is extensive literature on resilience to false-data injection at the sensor/estimator level [1]–[6], as well as anomaly detection methods that seek to identify unusual behaviors in data streams [7]–[9]. However, there has been less work done on secure state estimation in multi-agent, dynamic games, which we distinguish from the literature that uses game-theoretic models of attacks on control system inputs and measurements [10]–[12].

This paper considers a discrete-time linear system with two types of measurements: In addition to the usual linear measurement, a single quadratic measurement is also made. We assume that the linear measurement can be attacked (i.e., corrupted by an adversary) while the quadratic measurement cannot be attacked. Though our model does not reference a multi-agent game, it is motivated by a game-theoretic setting where the quadratic measurement corresponds to a utility-function value or reward that is received by the entire system. The goal of this paper is two-fold: To develop an observer for quadratic measurements, and to develop a statistical testing framework to detect attacks on the linear measurements.

*This material is based upon work supported by the National Science Foundation under Grant DGE-2125913.

The authors are with the department of Industrial Engineering and Operations Research, University of California, Berkeley, CA 94720, USA. {muyan-jiang, aaswani}@berkeley.edu

A. Collusion Detection in Multi-Agent Games

A closely related topic with increasing attention is detecting collusion in multi-agent games [13]–[18]. One set of approaches that has been proposed to detect collusion is the use of statistical tests [15], [17]. Another set of approaches leverage classical artificial intelligence (AI) [13], [14], [18]. For example, in large-scale team-based games, systems combining social networks with play metadata and unsupervised anomaly detection have been used to flag suspicious pairs [18]. In repeated-game scenarios, model-agnostic tests that retrain or simulate counterfactual strategies can reveal latent collusion by checking whether an agent’s policy becomes more exploitable under alternative assumptions [16]. Our work is related because it involves detecting undesired behavior in multi-agent games, but it differs in the type and model of undesired behavior.

B. Observer Design for Quadratic Measurements

Observer design for quadratic measurements is a less well-studied topic. One approach to observer design is to augment the state with derivatives or auxiliary variables of the quadratic output, which under certain conditions on the system convert the problem into an equivalent higher-dimensional linear one and enabling Kalman-like observers with convergence guarantees [19]. Related efforts analyze control for linear–quadratic output systems, including stabilizability criteria [20], and observability results for position estimation using only range or bearing data [21]. These works address systems with only quadratic outputs and no adversarial interference. By contrast, we design and apply a novel observer for quadratic measurements to a system with a vulnerable linear channel, enforcing consistency with the quadratic measurement to yield an attack-resilient observer.

C. Contributions and Outline

We make two main contributions in this paper. The first is that we develop a novel state observer for quadratic measurements. The second is that we develop a statistical test that uses the quadratic measurement to identify when the linear measurement is being attacked.

Section II presents the system model. Section III defines our novel observer for quadratic measurements, and performs a theoretical analysis. Section IV designs a statistical test for detecting attacks on the linear measurement. Section V reports numerical experiments on a pursuit–evasion game, demonstrating detection of a sensor attack and maintenance of estimation accuracy under an attack.

II. SYSTEM MODEL

This section presents the discrete-time linear system and its measurement model, and then it provides a game-theoretic interpretation of the quadratic measurements in this model.

A. Dynamics and Measurements

Consider a linear system $x_{k+1} = Ax_k + Bu_k + w_k$ with stochastic process noise $w_k \sim \mathcal{N}(0, Q)$, where $x_k \in \mathbb{R}^n$ is the state at time k , $u_k \in \mathbb{R}^m$ is the control input, $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$ are known system matrices, and $w_k \in \mathbb{R}^n$ is zero-mean Gaussian process noise with covariance $Q \succ 0$.

Assumption 1 (Invertibility): We make an assumption that the matrix $A \in \mathbb{R}^{n \times n}$ is invertible.

At each time step, two types of measurements are made: The linear measurements are $y_k = Cx_k + a_k + v_k$, with $v_k \sim \mathcal{N}(0, R)$, where $C \in \mathbb{R}^{p \times n}$ is the measurement matrix, $v_k \in \mathbb{R}^p$ is Gaussian measurement noise, and $a_k \in \mathbb{R}^p$ is an unknown attack vector that may corrupt y_k . The quadratic measurements are $z_k = x_k^\top V x_k$, with $V \in \mathbb{R}^{n \times n}$ and $V \succ 0$, and this measurement cannot be manipulated by the attacker.

B. Game-Theoretic Interpretation

Our motivation for studying the above model is the following game-theoretic interpretation: There are multiple agents subject to the dynamics, who pick inputs $u_k = g(x_k)$ to achieve a game-theoretic equilibrium (e.g., Nash equilibrium), and they have perfect knowledge of the $g(\cdot)$ equilibrium control policy. The linear measurements are susceptible to false-data injection [2], [3], [5], while the quadratic measurement is a game-theoretic, realized system utility or reward. The quadratic measurements may be a physical quantity (e.g., energy or Euclidean distance) obtained from local sensors and hence tamper-resistant [22].

III. OBSERVER FOR QUADRATIC MEASUREMENTS

Here, we present a novel observer for quadratic measurements. Then we theoretically analyze it.

A. Observer Design

Because the linear measurements are susceptible to attack, whereas the quadratic measurements are not, we use two observers: The first only uses linear measurements, and the second only uses quadratic measurements.

For linear measurements, we use a Kalman filter [23]:

Prediction:

$$\hat{x}_{k|k-1}^L = A\hat{x}_{k-1|k-1}^L + Bu_{k-1}$$

Update:

$$\hat{x}_{k|k}^L = \hat{x}_{k|k-1}^L + L_k(y_k - C\hat{x}_{k|k-1}^L),$$

where $P_{k|k-1}^L = AP_{k-1|k-1}^L A^\top + Q$ is predicted covariance, $P_{k|k}^L = (I - L_k C)P_{k|k-1}^L$ is updated covariance, and $L_k = P_{k|k-1}^L C^\top (CP_{k|k-1}^L C^\top + R)^{-1}$ is the Kalman gain.

For quadratic measurements, we propose an extended-Kalman-style observer, followed by a consistency projection:

Prediction:

$$\hat{x}_{k|k-1}^Q = A\hat{x}_{k-1|k-1}^Q + Bu_{k-1}$$

Extended Kalman Filter (EKF)-Like Correction:

$$\tilde{x}_{k|k} = \hat{x}_{k|k-1}^Q + K_k(z_k - (\hat{x}_{k|k-1}^Q)^\top V \hat{x}_{k|k-1}^Q),$$

where $H_k = (2V\hat{x}_{k|k-1}^Q)^\top$, $P_{k|k}^Q = (I - K_k H_k)P_{k|k-1}^Q$, $K_k = P_{k|k-1}^Q H_k^\top (H_k P_{k|k-1}^Q H_k^\top + \eta)^{-1}$, $\eta > 0$ regularizes the gain, and $P_{k|k-1}^Q = AP_{k-1|k-1}^Q A^\top + Q$.

Constrained Projection:

$$\hat{x}_{k|k}^Q = \arg \min_{x \in \mathcal{F}_k} \|x - \tilde{x}_{k|k}\|_{P_{k|k}^Q}^2,$$

where $\mathcal{F}_k = \bigcap_{i=0}^N \{x : |H_{k-i}(A^{-i}x - \hat{x}_{k-i|k-i}^Q) - \tilde{z}_{k-i}| \leq \delta_{k,i}(x)\}$, $\tilde{z}_{k-i} = z_{k-i} - (\hat{x}_{k-i|k-i}^Q)^\top V \hat{x}_{k-i|k-i}^Q$, $\delta_{k,i}(x) = \zeta + L\|A^{-i}x - \hat{x}_{k-i|k-i}^Q\|^2$, and $L = \|V\|_2$. The EKF-like correction step treats the quadratic measurement $z_k = x_k^\top V x_k$ as a nonlinear observation $h(x) = x^\top V x$. We linearize h around the prior $\hat{x}_{k|k-1}^Q$ via its Jacobian $H_k = (2V\hat{x}_{k|k-1}^Q)^\top$, and then apply a standard Kalman-style update with gain K_k .

While this captures the local curvature of the quadratic sensor, it can drift when the linearization is poor. To counteract this, we project the corrected estimate $\tilde{x}_{k|k}$ onto the feasible set $\mathcal{F}_k$. This set is defined by linearized measurement constraints from the current and past N steps, with adaptive bounds $\delta_{k,i}(x)$ that account for the second-order (linearization) error. By solving the projection, it returns the closest point, under the covariance-weighted norm, to the unconstrained update, while remaining compatible with all secure quadratic measurements. This enhances robustness by anchoring the estimate to true system behavior, even in the presence of large innovations or attacked linear signals.

B. Theoretical Error Bound

Here, we analyze the noise-free case. Since the inputs u_k are assumed to be known, without loss of generality we analyze our observer for the system: $x_{k+1} = Ax_k$ and $z_k = x_k^\top V x_k$, where V is symmetric and positive definite.

We begin by noting that the absolute value constraint

$$|H_{k-i}(A^{-i}x - \hat{x}_{k-i|k-i}^Q) - \tilde{z}_{k-i}| \leq \delta_{k,i}(x),$$

with $\delta_{k,i}(x) = \zeta + L\|A^{-i}x - \hat{x}_{k-i|k-i}^Q\|^2$, is equivalent to two inequalities. For $i = 0, \dots, N$, define

$$\varphi_i^+(x) = H_{k-i}(A^{-i}x - \hat{x}_{k-i|k-i}^Q) - \tilde{z}_{k-i} - \delta_{k,i}(x),$$

$$\varphi_i^-(x) = -H_{k-i}(A^{-i}x - \hat{x}_{k-i|k-i}^Q) + \tilde{z}_{k-i} - \delta_{k,i}(x),$$

so $\varphi_i^\pm(x) \leq 0$ encodes the same constraint. Since H_{k-i} and $\tilde{z}_{k-i}$ are constants and $\delta_{k,i}(x)$ is quadratic, each $\varphi_i^\pm$ is C^2 .

Next we establish the prox-regularity of the feasible set $\mathcal{F}_k$ using the theory of amenable sets [24], by making some mild assumptions about constraint qualification.

Assumption 2 (Nondegeneracy): For $i = 0, \dots, N$ and $s \in \{+, -\}$, if $\varphi_i^s(\bar{x}) = 0$ then $\nabla \varphi_i^s(\bar{x}) \neq 0$. Equivalently, if $y \in N_{(-\infty, 0]}(\varphi_i^s(\bar{x}))$ and $-\nabla \varphi_i^s(\bar{x})^* y = 0$, then $y = 0$.

Assumption 3 (Aggregated Constraint Qualification): Define the stacked mapping

$$F(x) = [\varphi_0^+(x) \quad \varphi_0^-(x) \quad \cdots \quad \varphi_N^+(x) \quad \varphi_N^-(x)]^\top \in \mathbb{R}^{2(N+1)}. \quad (1)$$

and let $D = \prod_{j=1}^{2(N+1)} (-\infty, 0]$. We assume that for $\bar{x}$: if $y \in N_D(F(\bar{x}))$ and $\nabla F(\bar{x})^* y = 0$, then $y = 0$.

Unless stated otherwise, all results in this subsection hold under Assumptions 1–3. We can formally define our feasible set as $\mathcal{F}_k = \{x \in \mathbb{R}^n : F(x) \in D\}$. This formulation allows us to establish the main result:

Proposition 1 (Prox-Regularity via Stacked Amenability): Under Assumptions 1–3, the set $\mathcal{F}_k$ is strongly amenable at $\bar{x}$ and, by [24, Proposition 13.32], prox-regular at $\bar{x}$.

Proof: Since each $\varphi_i^\pm$ is C^2 (due to its affine-plus-quadratic structure) and A is invertible by Assumption 1, the mapping $F : \mathbb{R}^n \rightarrow \mathbb{R}^{2(N+1)}$ is C^2 . The set $D = (-\infty, 0]^{2(N+1)}$ is closed, convex, and polyhedral.

By [24, Definition 10.23(b)], the representation $\mathcal{F}_k = \{x \in \mathbb{R}^n : F(x) \in D\}$ establishes that $\mathcal{F}_k$ is strongly amenable at $\bar{x}$ provided the constraint qualification

if $y \in N_D(F(\bar{x}))$ and $\nabla F(\bar{x})^* y = 0$, then $y = 0$

holds. Assumption 2 ensures that each active constraint φ_i^s is nondegenerate (i.e., $\nabla \varphi_i^s(\bar{x}) \neq 0$), while Assumption 3 guarantees the aggregated constraint qualification for F .

Therefore, by [24, Proposition 13.32], the indicator function $\delta_{\mathcal{F}_k}$ is prox-regular and subdifferentially continuous at $\bar{x}$. Equivalently, the set $\mathcal{F}_k$ is prox-regular at $\bar{x}$. ■

Lemma 1 (Feasibility of State with Adaptive Bounds):

In the noise-free case, the true state x_k belongs to the feasible set $\mathcal{F}_k$ when using the adaptive bounds $\delta_{k,i}(x) = \zeta + L\|A^{-i}x - \hat{x}_{k-i|k-i}\|^2$ where $L = \|V\|_2$ and $\zeta = 0$ in the noise-free case.

Proof: For the true state x_k to be in $\mathcal{F}_k$, it must satisfy: $|\tilde{z}_{k-i} - H_{k-i}(A^{-i}x_k - \hat{x}_{k-i|k-i})| \leq L\|A^{-i}x_k - \hat{x}_{k-i|k-i}\|^2$. From system dynamics, $x_{k-i} = A^{-i}x_k$, so we need to verify $\tilde{z}_{k-i} - H_{k-i}(x_{k-i} - \hat{x}_{k-i|k-i}) \leq L\|x_{k-i} - \hat{x}_{k-i|k-i}\|^2$.

Let $e_{k-i|k-i} = x_{k-i} - \hat{x}_{k-i|k-i}$. Substituting $H_{k-i} = (2V\hat{x}_{k-i|k-i})^\top$ and $\tilde{z}_{k-i} = z_{k-i} - \hat{x}_{k-i|k-i}^\top V \hat{x}_{k-i|k-i}$ gives $\tilde{z}_{k-i} - 2\hat{x}_{k-i|k-i}^\top V e_{k-i|k-i} \leq L\|e_{k-i|k-i}\|^2$. With $z_{k-i} = x_{k-i}^\top V x_{k-i}$ (noise-free case): $(2\hat{x}_{k-i|k-i}^\top V e_{k-i|k-i} + e_{k-i|k-i}^\top V e_{k-i|k-i}) - 2\hat{x}_{k-i|k-i}^\top V e_{k-i|k-i} = e_{k-i|k-i}^\top V e_{k-i|k-i} \leq \|V\|_2\|e_{k-i|k-i}\|^2 = L\|e_{k-i|k-i}\|^2$. So state x_k satisfies all constraints and belongs to $\mathcal{F}_k$. ■

Lemma 2 (Cross-Error Term Inequality): Define the pre-projection error $\tilde{e}_{k+1|k+1} \triangleq \tilde{x}_{k+1|k+1} - x_{k+1}$ and the projection error $e_{k+1}^{obj} \triangleq \tilde{x}_{k+1|k+1} - \hat{x}_{k+1|k+1}$. Then, under prox-regularity of $\mathcal{F}_{k+1}$, for any $x_{k+1} \in \mathcal{F}_{k+1}$ (in particular for the true state) $\tilde{e}_{k+1|k+1}^\top P_{k+1|k+1}^{-1} e_{k+1}^{obj} \geq \|e_{k+1}^{obj}\|_{P_{k+1|k+1}^{-1}}^2$.

Proof: Since $\hat{x}_{k+1|k+1}$ is a local minimizer of

$$\hat{x}_{k+1|k+1} = \arg \min_{x \in \mathcal{F}_{k+1}} \|x - \tilde{x}_{k+1|k+1}\|_{P_{k+1|k+1}^{-1}}^2,$$

The first-order necessary optimality condition for constrained optimization requires $-\nabla f(\hat{x}_{k+1|k+1}) \in N_{\mathcal{F}_{k+1}}(\hat{x}_{k+1|k+1})$, where $N_{\mathcal{F}_{k+1}}(\hat{x}_{k+1|k+1})$ is the proximal normal cone to $\mathcal{F}_{k+1}$ at $\hat{x}_{k+1|k+1}$, and $\nabla f(x) = 2P_{k+1|k+1}^{-1}(x - \tilde{x}_{k+1|k+1})$. So $2P_{k+1|k+1}^{-1}(\tilde{x}_{k+1|k+1} - \hat{x}_{k+1|k+1}) \in N_{\mathcal{F}_{k+1}}(\hat{x}_{k+1|k+1})$.

A key property of proximal normal cones for prox-regular sets is that for any $v \in N_{\mathcal{F}_{k+1}}(\hat{x}_{k+1|k+1})$ and any feasible

point $x \in \mathcal{F}_{k+1}$, we have $(x - \hat{x}_{k+1|k+1})^\top v \leq 0$ [24]. Applying this to our case with $v = 2P_{k+1|k+1}^{-1}(\tilde{x}_{k+1|k+1} - \hat{x}_{k+1|k+1}) = 2P_{k+1|k+1}^{-1}e_{k+1}^{obj}$ and $x = x_{k+1}$, we get $(x_{k+1} - \hat{x}_{k+1|k+1})^\top \cdot 2P_{k+1|k+1}^{-1}e_{k+1}^{obj} \leq 0$. Since $e_{k+1}^{obj} = \tilde{x}_{k+1|k+1} - \hat{x}_{k+1|k+1}$, substituting gives $x_{k+1} - \hat{x}_{k+1|k+1} = (x_{k+1} - \tilde{x}_{k+1|k+1}) + (\tilde{x}_{k+1|k+1} - \hat{x}_{k+1|k+1}) = -\tilde{e}_{k+1|k+1} + e_{k+1}^{obj}$, which implies $(-\tilde{e}_{k+1|k+1} + e_{k+1}^{obj})^\top P_{k+1|k+1}^{-1}(-e_{k+1}^{obj}) \geq 0$. Expanding this gives that we have $\tilde{e}_{k+1|k+1}^\top P_{k+1|k+1}^{-1}e_{k+1}^{obj} - (e_{k+1}^{obj})^\top P_{k+1|k+1}^{-1}e_{k+1}^{obj} \geq 0$, which implies that we have $\tilde{e}_{k+1|k+1}^\top P_{k+1|k+1}^{-1}e_{k+1}^{obj} \geq \|e_{k+1}^{obj}\|_{P_{k+1|k+1}^{-1}}^2$. ■

Theorem 1 (Projection Error Bound): Under the prox-regularity of $\mathcal{F}_{k+1}$, the projection step guarantees the post-projection error is bounded by the pre-projection error in the weighted norm $\|e_{k+1|k+1}\|_{P_{k+1|k+1}^{-1}}^2 \leq \|\tilde{e}_{k+1|k+1}\|_{P_{k+1|k+1}^{-1}}^2$, where the post-projection error is defined as $e_{k+1|k+1} \triangleq x_{k+1} - \hat{x}_{k+1|k+1}$.

Proof: We have $e_{k+1|k+1} = x_{k+1} - \hat{x}_{k+1|k+1} = (x_{k+1} - \tilde{x}_{k+1|k+1}) + (\tilde{x}_{k+1|k+1} - \hat{x}_{k+1|k+1}) = -\tilde{e}_{k+1|k+1} + e_{k+1}^{obj}$. Thus, $\|e_{k+1|k+1}\|_{P_{k+1|k+1}^{-1}}^2 = \|\tilde{e}_{k+1|k+1} - e_{k+1}^{obj}\|_{P_{k+1|k+1}^{-1}}^2 = \|\tilde{e}_{k+1|k+1}\|_{P_{k+1|k+1}^{-1}}^2 + \|e_{k+1}^{obj}\|_{P_{k+1|k+1}^{-1}}^2 - 2\tilde{e}_{k+1|k+1}^\top P_{k+1|k+1}^{-1}e_{k+1}^{obj}$. By Lemma 2, $\tilde{e}_{k+1|k+1}^\top P_{k+1|k+1}^{-1}e_{k+1}^{obj} \geq \|e_{k+1}^{obj}\|_{P_{k+1|k+1}^{-1}}^2$, and so we have $\|e_{k+1|k+1}\|_{P_{k+1|k+1}^{-1}}^2 \leq \|\tilde{e}_{k+1|k+1}\|_{P_{k+1|k+1}^{-1}}^2 - \|e_{k+1}^{obj}\|_{P_{k+1|k+1}^{-1}}^2 \leq \|\tilde{e}_{k+1|k+1}\|_{P_{k+1|k+1}^{-1}}^2$. ■

IV. STATISTICAL TEST FOR ATTACK DETECTION

This section develops a statistical test to detect attacks on the linear measurements, using the unattackable quadratic measurements. More formally, suppose the null hypothesis (H_0) is that the state estimate distributions of the two observers coincide. Conversely, under the alternative hypothesis (H_1), an adversarial attack injects a perturbation signal a_k , causing the linear observer's state estimate distribution to deviate significantly from that of the quadratic observer.

A key challenge is that observer estimates are temporally dependent, while standard permutation tests assume independence. Hence we use the wild bootstrap maximum mean discrepancy (MMD) test [25], which is designed for time-dependent data such as state trajectories. Let $X_k^L = \{\hat{x}_{1|1}^L, \dots, \hat{x}_{k|k}^L\}$ and $X_k^Q = \{\hat{x}_{1|1}^Q, \dots, \hat{x}_{k|k}^Q\}$ represent historical state estimates from the linear and quadratic observers respectively. To quantify the discrepancy between these empirical distributions, we adopt the MMD equipped with an RBF kernel $\phi(x, y) = \exp(-\|x - y\|^2 / 2\sigma^2)$. The empirical squared MMD is $\text{MMD}^2(X_k^L, X_k^Q) = \frac{1}{k^2} \sum_{i,j=1}^k (\phi(\hat{x}_{i|1}^L, \hat{x}_{j|1}^L) + \phi(\hat{x}_{i|1}^Q, \hat{x}_{j|1}^Q) - 2\phi(\hat{x}_{i|1}^L, \hat{x}_{j|1}^Q))$.

To assess statistical significance while preserving temporal dependence, we implement the wild bootstrap approach as follows. First, combine the estimates into a single set $Z_k = \{\hat{x}_{1|1}^L, \dots, \hat{x}_{k|k}^L, \hat{x}_{1|1}^Q, \dots, \hat{x}_{k|k}^Q\}$ with $2k$ total observations. Construct the kernel matrix $K \in \mathbb{R}^{2k \times 2k}$ with entries $K_{ij} =$

$\phi(Z_i, Z_j)$, and center it using the centering matrix $H = I_{2k} - \frac{1}{2k} \mathbf{1}_{2k} \mathbf{1}_{2k}^\top$ to obtain the centered kernel matrix $\tilde{K} = HKH$.

Next, define random perturbation variables $\{v_i\}_{i=1}^{2k}$, independently drawn from a symmetric distribution with mean zero and unit variance (e.g., the Rademacher distribution). Using these, construct the wild bootstrap kernel matrix via element-wise perturbation: $\tilde{K}_{ij}^v = v_i v_j \tilde{K}_{ij}$. Then, compute the bootstrap MMD statistic for each realization as $\text{MMD}_v = \frac{1}{2k} \sum_{i,j=1}^{2k} \tilde{K}_{ij}^v$. This sum is a degenerate V-statistic, and it mimics the null distribution of MMD^2 under dependence. Repeating this bootstrap procedure B times generates a distribution of bootstrap statistics: $\{\text{MMD}_v^{(1)}, \dots, \text{MMD}_v^{(B)}\}$, from which we derive a critical threshold γ_α at significance level α (by, for example, using the $(1 - \alpha)$ -quantile).

Finally, the decision rule for attack detection becomes $\text{MMD}^2(X_k^L, X_k^Q) \stackrel{H_1}{\underset{H_0}{\gtrless}} \gamma_\alpha$. By adopting this wild bootstrap strategy, the statistical test accommodates temporal correlations within the state estimate sequences, ensuring valid inference even in the presence of inherent data dependence.

V. NUMERICAL EXPERIMENTS

We conduct numerical experiments on a two-agent pursuit-evasion game governed by double integrator dynamics. Although our theoretical results assumed a noise-free regime, we include moderate Gaussian noise to demonstrate robustness beyond theoretical guarantees.

A. Experimental Setup

We consider a planar two-agent system with state vector $x_k \in \mathbb{R}^8$ at discrete time k , given by $x_k = [p_A, v_A, p_B, v_B]^\top$, where $p_A, p_B \in \mathbb{R}^2$ denote the positions and $v_A, v_B \in \mathbb{R}^2$ the velocities of the evader (Agent A) and pursuer (Agent B), respectively. The system evolves according to the discrete-time double integrator model $x_{k+1} = Ax_k + Bu_k + w_k$, where $A \in \mathbb{R}^{8 \times 8}$ and $B \in \mathbb{R}^{8 \times 4}$ are the state transition and input matrices, $u_k = [u_A, u_B]^\top$ is the control input, and w_k is zero-mean Gaussian process noise. The measurement model includes two channels: a vulnerable linear measurement $y_k = Cx_k + a_k + v_k$, where C extracts the positions of both agents, $v_k \sim \mathcal{N}(0, R)$ is Gaussian noise, and a_k is an adversarial attack vector; and a secure quadratic measurement $z_k = x_k^\top V x_k$, where $V \in \mathbb{R}^{8 \times 8}$ encodes the relative Euclidean distance.

The simulation parameters are as follows: time step $\Delta t = 0.1$ s, simulation horizon 20 steps, and process/measurements noise standard deviations all set to 0.005. To assess robustness, we use randomized initial conditions drawn from Gaussian neighborhoods: the evader position $p_A(0)$ is sampled around $(0, 0)$ with standard deviation 0.5 in each axis, and the pursuer position $p_B(0)$ is sampled around $(2, 2)$ with standard deviation 1.5. Initial velocities have random directions (uniform over the unit circle) and magnitudes drawn from $\mathcal{N}(0.5, 0.05^2)$ for the evader and $\mathcal{N}(0.2, 0.05^2)$ for the pursuer, truncated below at 0.1 m/s.

We run $M = 100$ independent trials with the above randomized initializations. For each time step, we aggregate

metrics across runs and report the mean together with the standard error (SE). The attack (defined in Sec. V-C and below) is injected at discrete time $k = 10$ with magnitude $\beta = 7.0$ along the relative position direction.

B. Control Policies

For discrete-time linear dynamics, optimal policies can be computed via Hamilton–Jacobi–Bellman–Isaacs (HJBI) formulations [26]–[29], and capture conditions under full observability and sufficient control authority are well established [28], [30]. For simplicity, we use heuristic control policies inspired by reachability-based strategies [31] and observer-based estimation frameworks [32], [33]. Control inputs are constrained component-wise by a saturation operator $[u]_{a_{\max}} \triangleq \max\{-a_{\max}, \min(u, a_{\max})\}$, $a_{\max} = 3 \text{ m/s}^2$. This prevents physically unrealistic actuator demands.

Pursuer (Agent B, Leader): Agent B uses perfect state knowledge to pursue an intercept point computed via one-step extrapolation with short-horizon interception timing:

- 1) **Evader Prediction:** Predict the evader’s next position: $\tilde{p}_A(k+1) = p_A(k) + v_A(k)\Delta t$.
- 2) **Intercept Calculation:** Let $d_k \triangleq \|p_A(k) - p_B(k)\|_2$, $r_k \triangleq \|v_A(k)\|_2$. If the evader is moving significantly (i.e., $\|v_A(k)\|_2 > 0.1 \text{ m/s}$), determine intercept time t^* by solving: $\|d_k + t(v_A(k) - v_B(k))\|_2^2 = (0.1 r_k)^2$, and set the intercept point as: $p_I = p_A(k) + v_A(k)t^*$. Otherwise, default to the simple extrapolation: $p_I = \tilde{p}_A(k+1)$.
- 3) **Desired Velocity:** The desired pursuit velocity combines range-dependent speed and near-range velocity matching: $v_B^{\text{des}}(k) = s(r_k) \cdot (p_I - p_B(k))/\|p_I - p_B(k)\|_2 + \beta(r_k)v_A(k)$, with: $s(r_k) = v_{\max, B} = 2.5$ if $r_k > 2$ and $s(r_k) = v_{\max, B} \cdot (0.5 + 0.25r_k)$ if $r_k \leq 2$ and $\beta(r_k) = 0.5$ if $r_k < 1$ and $\beta(r_k) = 0$ otherwise. The pursuer thus aggressively pursues at larger distances but smoothly transitions to cautious, velocity-matched intercept as the range closes, inspired by practical intercept strategies validated in [31].
- 4) **Control Law:** The control input for Agent B is computed as $u_B(k) = [(v_B^{\text{des}}(k) - v_B(k))/\Delta t]_{a_{\max}}$.

Evader (Agent A, Follower): The evader relies exclusively on the observer estimate $\hat{x}_k$ and strategically evades by forecasting the pursuer’s short-term motion:

- 1) **Pursuer Prediction:** Predict the pursuer’s next position from the estimate: $\tilde{p}_B(k+1) = \hat{p}_B(k) + \hat{v}_B(k)\Delta t$.
- 2) **Escape Direction:** Compute the escape direction from the predicted pursuer position: $e_k = \hat{p}_A(k) - \tilde{p}_B(k+1)$.
- 3) **Desired Velocity:** Maximize distance along the escape vector and add a minor velocity-matching perturbation to introduce unpredictability [32], [33] at longer distances: $v_A^{\text{des}}(k) = v_{\max, A} \cdot e_k/\|e_k\|_2 + \gamma(\hat{r}_k)\hat{v}_B(k)$, with $v_{\max, A} = 1.5$, $\hat{r}_k = \|\hat{p}_A(k) - \hat{p}_B(k)\|_2$, and $\gamma(\hat{r}_k) = 0.2$ if $\hat{r}_k > 2$, otherwise $\gamma(\hat{r}_k) = 0$.
- 4) **Control Law:** The control input for Agent A is similarly computed using estimated states: $u_A(k) = [(v_A^{\text{des}}(k) - \hat{v}_A(k))/\Delta t]_{a_{\max}}$.




Fig. 1. Representative trial from the repeated experiments. Left: linear observer. Right: quadratic observer. True trajectories are shown for the evader (blue) and pursuer (red). The estimated pursuer trajectory is overlaid with hollow circles: green for pre-attack samples and orange for post-attack samples. Yellow “X” marks the attack onset. Faint lines trace the motion path and arrows indicate instantaneous velocity.

The realized control inputs $u_k = [u_A(k); u_B(k)]$ are assumed to be known by both observers.

C. Attack Scenario

To evaluate detection and estimation robustness, we inject a *relative position attack* on the linear channel at time $k = 10$. The attack vector is constructed as $a_k = \beta \frac{p_B - p_A}{\|p_B - p_A\|}$, where $\beta = 7.0$ is the attack magnitude. We interpret β as a distance bias magnitude in meters, injected along the relative position vector. This attack biases the perceived position of the pursuer, misleading the vulnerable observer.

D. Experimental Results

1) *Trajectory Analysis*: Fig. 1 shows a single representative trial drawn from the repeated-experiment protocol with randomized initial positions and velocities. Under the linear Kalman observer (left), the estimated pursuer trajectory (hollow orange circles, post-attack) departs from the red ground-truth path immediately after the attack marker (yellow “X”). The drift appears as a systematic, directionally consistent bias that grows along the motion direction, yielding a spurious “phantom” pursuer that advances more slowly and farther from truth. In contrast, the quadratic observer (right) remains well aligned with the true pursuer trajectory both before and after the attack; the hollow green (pre-attack) and orange (post-attack) estimates closely overlay the red curve. Comparing the two panels over the same time horizon, the pursuer under the vulnerable linear observer appears to close the gap to the evader more than under the quadratic observer. Fig. 2 summarizes the mean squared error (MSE) between true states and observer estimates across the repeated runs. Before the attack, both observers achieve comparable accuracy. After attack, the MSE of the linear observer increases markedly, whereas the quadratic observer maintains a low error by using the secure quadratic measurement.

2) *Attack Detection*: We use an RBF kernel (width via the median heuristic), and a wild bootstrap with Rademacher




Fig. 2. Observer MSE over time aggregated across $M = 100$ runs. Red: linear observer; blue: quadratic observer. Solid lines: mean MSE; shaded regions: $\pm$ SE. Vertical dashed line indicates attack onset.

multipliers, $B = 500$, $\alpha = 0.05$. Online evaluation uses a sliding window W equal to the pre-attack horizon; we declare detection at time k if $\text{MMD}_k^2 > \hat{\gamma}_{\alpha,k}$.

Fig. 3 reports the *aggregated* wild bootstrap MMD statistic across $M = 100$ runs (mean $\pm$ SE) together with the corresponding mean critical value (dashed). Prior to the attack, the statistic remains below the threshold with no false positives on average. At the attack onset (vertical line), the mean MMD crosses the critical value with no delay, and the margin continues to widen thereafter, indicating a persistent distributional divergence between the drifted linear-observer trajectory and the stable quadratic-observer trajectory.

The results demonstrate that, with accurate initialization and low noise, the quadratic observer maintains robust state estimation in the presence of adversarial attacks, while the linear Kalman observer is significantly compromised. The MMD-based test provides prompt and reliable attack detection. These findings validate the theory and highlight the practical utility of the proposed approach for resilient state estimation and attack detection in dynamic games.




Fig. 3. Attack detection via wild bootstrap MMD over $M = 100$ repeated experiments. Green: mean MMD test statistic; orange: mean critical value. Solid lines: mean; shaded bands: $\pm$ SE. Vertical dotted line marks attack onset; red shaded region denotes the attack period.

VI. CONCLUSION

This work presented a robust framework for detecting adversarial sensor attacks in linear dynamical systems by combining a novel quadratic observer with a wild bootstrap MMD test. The quadratic observer leverages secure quadratic measurements to maintain reliable state estimates, while the wild bootstrap test detects distributional shifts under temporal dependence. Our theoretical analysis established error-monotonicity and prox-regularity properties of the proposed observer, and numerical experiments on a pursuit–evasion game demonstrated accurate estimation and prompt attack detection. Future work includes scaling the framework to larger multi-agent systems, incorporating adaptive thresholds for online testing, and extending the approach to nonlinear dynamics and broader classes of adversarial strategies.

REFERENCES

- [1] Y. Mo and B. Sinopoli, “Secure estimation in the presence of integrity attacks,” *IEEE Transactions on Automatic Control*, vol. 60, no. 4, pp. 1145–1151, 2015.
- [2] F. Pasqualetti, F. Dörfler, and F. Bullo, “Attack detection and identification in cyber-physical systems,” *IEEE Transactions on Automatic Control*, vol. 58, no. 11, pp. 2715–2729, 2013.
- [3] F. H. Fawzi, P. Tabuada, and S. Diggavi, “Secure estimation and control for cyber-physical systems under adversarial attacks,” *IEEE Transactions on Automatic Control*, vol. 59, no. 6, pp. 1454–1467, 2014.
- [4] P. Hespanhol, M. Porter, R. Vasudevan, and A. Aswani, “Dynamic watermarking for general lti systems,” in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 2017, pp. 1834–1839.
- [5] M. Pajic, I. Lee, and G. J. Pappas, “Attack-resilient state estimation for noisy dynamical systems,” *IEEE Transactions on Control of Network Systems*, vol. 4, no. 1, pp. 82–92, 2017.
- [6] T. Yang, C. Murguia, M. Kuijper, and D. Nešić, “A multi-observer framework for nonlinear systems under sensor attacks,” *Automatica*, vol. 119, p. 109043, 2020.
- [7] M. H. Bhuyan, D. K. Bhattacharyya, and J. K. Kalita, “Network anomaly detection: methods, systems and tools,” *Ieee communications surveys & tutorials*, vol. 16, no. 1, pp. 303–336, 2013.
- [8] A. Blázquez-García, A. Conde, U. Mori, and J. A. Lozano, “A review on outlier/anomaly detection in time series data,” *ACM computing surveys (CSUR)*, vol. 54, no. 3, pp. 1–33, 2021.
- [9] G. Pang, C. Shen, L. Cao, and A. V. D. Hengel, “Deep learning for anomaly detection: A review,” *ACM computing surveys (CSUR)*, vol. 54, no. 2, pp. 1–38, 2021.

- [10] H. J. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram, “Resilient asymptotic consensus in robust networks,” *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 4, pp. 766–781, 2013.
- [11] Q. Zhu and T. Başar, “Game-theoretic methods for robustness, security, and resilience of cyber-physical control systems,” *IEEE Control Systems Magazine*, vol. 35, no. 1, pp. 46–65, 2015.
- [12] J. Pawlick, E. Colbert, and Q. Zhu, “A game-theoretic taxonomy and survey of defensive deception for cybersecurity and privacy,” *ACM Computing Surveys*, vol. 52, no. 4, pp. 82:1–82:28, 2019.
- [13] G. K. Palshikar and M. M. Apte, “Collusion set detection using graph clustering,” *Data Mining and Knowledge Discovery*, vol. 15, no. 3, pp. 279–298, 2007.
- [14] P. Mazrooei, C. Archibald, and M. Bowling, “Automating collusion detection in sequential games,” in *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 2013, pp. 675–681.
- [15] P. Hespanhol and A. Aswani, “Hypothesis testing approach to detecting collusion in competitive environments,” in *Proceedings of the 13th EAI International Conference on Performance Evaluation Methodologies and Tools*, 2020, pp. 35–40.
- [16] M. Courthoud, “Algorithmic collusion detection,” 2021, model-free test for detecting algorithmic collusion from observational or retraining-based evidence.
- [17] T. Bonjour, V. Aggarwal, and B. Bhargava, “Information-theoretic approach to detect collusion in multi-agent games,” in *Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence (UAI)*, ser. Proceedings of Machine Learning Research, vol. 180, 2022, pp. 223–232.
- [18] L. Greige, F. D. M. Silva, M. Trotter, C. Lawrence, P. Chin, and D. K. Varadarajan, “Collusion detection in team-based multiplayer games,” *arXiv*, vol. 2203.05121, 2022. [Online]. Available: <https://arxiv.org/abs/2203.05121>
- [19] D. Theodosis, S. Berkane, and D. V. Dimarogonas, “State estimation for a class of linear systems with quadratic output,” in *Proceedings of the 24th International Symposium on Mathematical Theory of Networks and Systems (MTNS)*, Cambridge, UK, 2021, pp. 261–266.
- [20] J.-M. Montenbruck, S. Zeng, and F. Allgöwer, “Linear systems with quadratic outputs,” in *Proceedings of the 2017 American Control Conference (ACC)*, 2017, pp. 1030–1034.
- [21] T. Hamel and C. Samson, “Position estimation from direction or range measurements,” *Automatica*, vol. 82, pp. 137–144, 2017.
- [22] S. Capkun, M. Hamdi, and J.-P. Hubaux, “Gps-free positioning in mobile ad-hoc networks,” in *Proceedings of the 34th Annual Hawaii International Conference on System Sciences*, 2001, pp. 10 pp.–.
- [23] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Transactions of the ASME — Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [24] R. T. Rockafellar and R. J.-B. Wets, *Variational analysis*. Springer Science & Business Media, 2009, vol. 317.
- [25] K. P. Chwialkowski, D. Sejdinovic, and A. Gretton, “A wild bootstrap for degenerate kernel tests,” *Advances in Neural Information Processing Systems*, vol. 27, 2014.
- [26] R. Isaacs, *Differential Games*. John Wiley & Sons, 1965.
- [27] M. Bardi and I. Capuzzo-Dolcetta, *Optimal Control and Viscosity Solutions of Hamilton–Jacobi–Bellman Equations*. Birkhäuser, 1997.
- [28] M. Falcone, “Numerical methods for differential games based on partial differential equations,” *International Game Theory Review*, vol. 3, no. 2, pp. 123–148, 1994.
- [29] C. J. Tomlin, J. Lygeros, and S. S. Sastry, “A game theoretic approach to controller design for hybrid systems,” *Proceedings of the IEEE*, vol. 88, no. 7, pp. 949–970, 2000.
- [30] Z. Li, P. Gao, and X. Dong, “Optimal capture strategies for 1-vs-1 discrete-time double integrator pursuit–evasion games,” *IEEE Transactions on Automatic Control*, 2024, to appear.
- [31] T. H. Chung and T. Furukawa, “Time-optimal control in pursuit–evasion with reachability constraints,” *IEEE Transactions on Robotics*, vol. 24, no. 5, pp. 873–880, 2008.
- [32] K. Horák and B. Bošanský, “Solving partially observable pursuit–evasion games using value iteration,” *Journal of Artificial Intelligence Research*, vol. 58, pp. 13–20, 2017.
- [33] Y. Oshman and P. Davidson, “Optimization of observer trajectories for bearings-only target localization,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 36, no. 4, pp. 1456–1469, 2000.