

Resilient Parameter-Invariant Control with Application to Vehicle Cruise Control

James Weimer, Nicola Bezzo, Miroslav Pajic,
George J. Pappas, Oleg Sokolsky, and Insup Lee

School of Engineering and Applied Sciences
University of Pennsylvania
Philadelphia, PA 19104, USA
{weimerj,nicbezzo,pajic,pappasg}@seas.upenn.edu
{sokolsky,lee}@cis.upenn.edu

Abstract. This work addresses the general problem of resilient control of unknown stochastic linear time-invariant (LTI) systems in the presence of sensor attacks. Motivated by a vehicle cruise control application, this work considers a first order system with multiple measurements, of which a bounded subset may be corrupted. A frequency-domain-designed resilient parameter-invariant controller is introduced that simultaneously minimizes the effect of corrupted sensors, while maintaining a desired closed-loop performance, invariant to unknown model parameters. Simulated results illustrate that the resilient parameter-invariant controller is capable of stabilizing unknown state disturbances and can perform state trajectory tracking.

Keywords: Secure Cyber-Physical Systems, Robust Control, Resilient Sensor Fusion

1 Introduction

Modern large-scale control systems are becoming more and more integrating into our daily lives. As the integration of smart devices in modern control systems increases, so does the need for defenses against potential attacks. Today, our dependence on integrated controllers automates everything from inter-home appliances to nation-wide power distribution, where the effect of unpredicted behaviors can range from a minor inconvenience of resetting a smart device to a regional blackout. Our dependence on these closed-loop automated systems requires that their performance be resilient to malicious behavior and robust to non-malicious behavior.

With respect to the vehicle cruise control system, non-malicious agents include environmental variables (gravity, wind speed, parts fatigue and failure, etc.) while malicious behavior can be introduced through, for example, sensor spoofing. Employing redundant measurements is a well established method of providing better estimates of control variables and model parameters; however, when attacked, a redundant measurement can be used as a means to destabilize a

control system. Thus, to ensure safe performance of the vehicle cruise controller requires securing the sensory data, while simultaneously designing controllers robust to unknown environmental parameters.

Literature review: The design of algorithms which are resilient against faults or unknown parameters has been addressed from many points of view, including fault detection [1], robust control [2], adaptive control [3], and more generally from estimation and hypothesis testing [4]. In general, these approaches address the issue of maximizing some performance measure with respect to known or bounded disturbances. In the context of security against malicious attacks, many of these approaches are not applicable because of their assumption that the attack is either known or bounded, with notable exceptions being approaches which ask for invariance to the unknown parameters [5]. The remainder of this literature review focuses on secure estimation/control and control of unknown systems, respectively.

Secure estimation and control system design in the presence of disturbances or attacks has received increasing research interest [6,7,8,9,10,11]. Most closely related to the work presented herein is [12], which addresses the secure estimation and control of linear deterministic systems under malicious sensor attacks. While the approach in [12] is shown to stabilize the systems under consideration, their approach requires full knowledge of the underlying system dynamics in order to secure the closed-loop system. When the underlying dynamics are unknown (and potentially stochastic) more robust detection and control algorithms are needed.

Control of unknown linear systems can be approached through adaptive control techniques [3], typically based on a Lyapunov stability requirement. Adaptive control is generally classified as either indirect (estimation of model parameters) or direct (estimation of the control sequence), where direct adaptive control is more robust than indirect adaptive control, since the parameter estimation may not be accurate over all frequencies [13,14,15]. However, adaptive control techniques ultimately require an estimation of either the unknown parameters or unknown control feedback gain, which dictates the performance of the closed-loop system. In contrast, in this work we wish to directly calculate the controller outputs for systems with unknown parameters, without requiring the intermediate estimation of the unknown parameters or feedback control gain, such that a closed-loop specification is satisfied.

Statement of contributions: Beyond the previous approaches, this work focuses on the specification-based design of resilient parameter-invariant controllers for stochastic linear time-invariant systems, with specific interest in vehicle cruise control systems. The primary technical contributions of this work are: (a) a resilient sensor fusion strategy for unknown attacks on noisy measurements; (b) a finite-horizon mean-stabilizing parameter-invariant controller; (c) a sensor fusion and controller co-design requirement that satisfies a stochastic Lyapunov criteria.

Structure of the paper: Section 2 formulates precisely the problem considered in this work. We introduce the resilient sensor fusion strategy and parameter-invariant controller in Section 3. A co-design requirement for satisfying the performance specification is provided in Section 4. Section 5 presents numerical evaluations of the resilient parameter-invariant controller in the presence of sensor attacks for vehicle cruise control. The concluding section provides discussion and proposes future extensions. For readability, the notation and constant definitions employed in this paper are included as an appendix.

2 Problem Formulation

This section introduces a resilient control problem for a system with unknown LTI-Gaussian dynamics and (potentially) corrupted measurements.

Specifically, we consider an LTI system with a single state that evolves according to

$$\begin{aligned} x(k+1) &= ax(k) + bu(k) + w(k) \\ \mathbf{y}(k) &= \mathbf{c}x(k) + \mathbf{v}(k) + \mathbf{d}(k) \end{aligned} \quad (1)$$

where:

- $x, u \in \mathbb{R}$, are the state and control input, respectively;
- $a, b \in \mathbb{R}$ are the state dynamic and control input gains;
- $\mathbf{y}, \mathbf{c}, \mathbf{d} \in \mathbb{R}^N$, are the measurements, state measurement gain, and corruption, respectively;
- $w \in \mathbb{R}$ and $\mathbf{v} \in \mathbb{R}^N$ are uncorrelated i.i.d. Gaussian process noise and measurement noise with central moments¹:

$$\begin{aligned} \mathbb{E}[w] &= \mu & \mathbb{E}[\mathbf{v}] &= \mathbf{0} \\ \mathbb{E}[(w - \mu)^2] &= \sigma & \mathbb{E}[(\mathbf{v}\mathbf{v}^\top)^2] &= \mathbf{I} \end{aligned}$$

At time k , the model information available in this work is summarized in the following assumption:

Assumption 1 - available information:

- the time-series measurements, $\mathbf{y} = [\mathbf{y}^\top(0), \dots, \mathbf{y}^\top(k)]^\top$;
- the time-series control inputs, $\mathbf{u} = [u(0), \dots, u(k)]^\top$;
- the variance of the process noise, σ ;
- the state measurement gain, \mathbf{c} ;
- the state dynamics, control gain, and process noise mean are constant.

¹ Without loss of generality, we assume the measurement noise is white and normalized to unit variance, where colored noise and non-unit variance white noise can be whitened by applying a normalizing pre-whitening filter.

In words, we assume that the measurements, measurement state gain, inputs, and noise covariances are known; however, the dynamics and process bias governing the evolution of the state are unknown, but constant. For completeness, we summarize the unavailable information in the following assumption²:

Assumption 2 - unavailable information:

- the state dynamics, a
- the control gain, b ;
- the process noise mean, μ ;
- the time-series measurement corruption, $\mathbf{d} = [\mathbf{d}^\top(0), \dots, \mathbf{d}^\top(k)]^\top$.

Although we assume the measurement corruption is unknown, we assume a maximum of M measurements are corrupted, as defined in the following assumption:

Assumption 3 - measurement corruption structure: At each time step, at most M measurements are corrupted, $\|\mathbf{d}(k)\|_0 \leq M$, where

$$M = \begin{cases} \frac{N}{2} - 1, & N \text{ even} \\ \frac{N-1}{2}, & N \text{ odd} \end{cases}$$

such that for

$$\{\mathbf{P}_d^\perp\} := \{\mathbf{Q}^\top \mathbf{Q} \mid \mathbf{Q} \in \{0, 1\}^{T \times N}, \mathbf{Q}\mathbf{Q}^\top = \mathbf{I}, T > N - M\}$$

the following is true:

$$\mathbf{F}\mathbf{d}(k) = \mathbf{0}, \quad \exists \mathbf{F} \in \{\mathbf{P}_d^\perp\}$$

Assumption 3 is consistent with the assumption in [12] and implies a maximum of M corrupted measurements since elements of $\mathbf{d}(k)$ which equal zero, imply no corruption is applied to the corresponding measurement. Under the assumptions introduced in this section, we wish to solve the following problem:

Problem 1 - stochastic boundedness: Given (1) and assumptions 1-3, show that

$$\mathbb{E} [\|x(k+1)\|^2] \leq \alpha^* \|x(k)\|^2, \quad \forall \|x(k)\|^2 \geq \eta \quad (2)$$

where $\eta \in \mathbb{R}$ is a desired state convergence threshold and $\alpha^* \in [0, 1]$ denotes a desired state convergence rate.

² The notation $\|\cdot\|_0$ denotes the zero-norm, while $\|\cdot\|$ denotes the two-norm.

The problem introduced in this section is addressed in the following section by introducing criteria for resilient sensor fusion and finite-horizon parameter-invariant control.

3 Main Contributions

By extensively employing the notation and definitions introduced in the appendix, the main contributions of this work are summarized in the following propositions:

Proposition 1 - resilient sensor fusion: Given (1), then

$$\hat{P}_d^\perp = \arg \min_{Q \in \{P_d^\perp\}} \frac{\left(\|P_{Q^c}^\perp Q y(k)\|^2 - \|Q\|_0 + 1 \right)^2}{2(\|Q\|_0 - 1)} \implies \mathbb{E} \left[\hat{P}_d^\perp d(k) \right] = \mathbf{0}$$

Proposition 2 - finite-horizon parameter-invariant control: Given (1), $\alpha \in [0, 1]$, positive integers $k_0, k, \kappa \in \mathbb{N}$ satisfying

$$0 < k - k_0 \leq \kappa \leq 2(k - k_0) - 4,$$

a $(\kappa - k_0 + 1)$ -point DFT matrix, \mathbf{V} , and

$$\mathbf{s}_{n-k_0} = \begin{cases} \mathbb{E}[x(n)|\mathbf{y}(n)], & n \in \{k_0, \dots, k-1\} \\ \alpha^{n-k} \mathbb{E}[x(k)|\mathbf{y}(k)], & n \in \{k, \dots, \kappa\} \end{cases}$$

$$\mathbf{H} = [\mathbf{A}\mathbf{V}\mathbf{s}, \quad \mathbf{V}\mathbf{s}, \quad \sum_{n=k_0}^{\kappa} \mathbf{V}_n, \quad \mathbf{V}_0]$$

then, assuming $\mathbf{u}_{k_0:k-1} \neq \mathbf{0}$,

$$P_H^\perp \mathbf{V} \mathbf{u}_{k_0:\kappa-1} = \mathbf{0} \implies \mathbb{E}[x(t)|\mathbf{y}_{k_0:k}] = \alpha^{t-k} \mathbb{E}[x(k)|\mathbf{y}(k)] \quad \forall k \leq t \leq \kappa$$

The following subsections discuss proposition 1 and proposition 2, respectively. The propositions discussed in this section are utilized in the following section to design a resilient model-invariant controller that satisfies the stochastic boundedness constraint in problem 1.

3.1 Resilient Sensor Fusion

Designing a stabilizing controller for problem 1 requires information feedback through sensor measurements. When corrupted by an unmodeled attacker, the

sensor contains no information with respect to the system state. Moreover, inclusion of corrupted measurements in state estimation infects the healthy information provided by the uncorrupted sensors. The purpose of resilient sensor fusion is to identify a set of sensors that are expected to be unaffected by the measurement corruption vector, $\mathbf{d}(k)$, (i.e. reside in the null space of $\mathbf{d}(k)$) and to generate a minimum mean squared estimate of the state, $x(k)$. From assumption 3, and consistent with the standard assumptions in the related work [12], we assume that there are at least $M + 1$ sensors which are uncorrupted. However, unlike the previous work, this problem considers uncorrupted measurements that are inherently noisy and that the specific sensors which are corrupted can change at each time-step. Coupled with the fact that the underlying dynamics are unknown (by assumption 2), the resilient sensor fusion strategy at time k is limited to using only sensor measurements provided at time k , and no prior information on the state (as prior information on the state will propagate previously corrupted information). In this subsection, we develop a resilient sensor fusion strategy that estimates the null space of the corruption, denoted as $\hat{\mathbf{P}}_d^\perp \in \{\mathbf{P}_d^\perp\}$, invariant to the value of the state and potential corruption, which is employed to generate a minimum mean squared error (MMSE) state estimate. Specifically, the design of the resilient sensor fusion strategy is organized into three steps:

1. formulate a test statistic;
2. test for the *best* corruption null space;
3. generate a MMSE state estimate.

These steps are respectively addressed in the remainder of this subsection.

Test Statistic: To determine the corrupted sensors requires identifying a test statistic that is invariant to the unknown system state which can be tested. Assuming a candidate corruption null space, $\mathbf{Q} \in \{\mathbf{P}_d^\perp\}$, and consistent with optimal signal detection theory [5], we write a maximally invariant statistic for $\mathbf{d}(k)$, invariant to the unknown state, $x(k)$, as

$$\mathbf{P}_{\mathbf{Q}^c}^\perp \mathbf{Q} \mathbf{y}(k) = \mathbf{P}_{\mathbf{Q}^c}^\perp \mathbf{Q} (\mathbf{d}(k) + \mathbf{v}(k)). \quad (3)$$

The maximally invariant statistic is a statistic of the measurements which contains all the information with respect to $\mathbf{Q} \mathbf{d}(k)$ that is invariant to the unknown state, $x(k)$. It is known that there exists no uniformly most powerful (UMP) test for detecting an unknown vectored signal [5], where the UMP test is considered optimal in hypothesis testing theory. Following well established practices employed in bad-data detection [5], we utilize the inner product of the maximally invariant statistic in (3),

$$\|\mathbf{P}_{\mathbf{Q}^c}^\perp \mathbf{Q} \mathbf{y}(k)\|^2 = \|\mathbf{P}_{\mathbf{Q}^c}^\perp \mathbf{Q} (\mathbf{d}(k) + \mathbf{v}(k))\|^2, \quad (4)$$

as a test statistic, in place of the maximally invariant statistic.

Corruption Null Space Test: To determine the corruption null space, we seek to identify the candidate \mathbf{Q} for the test statistic that *best fits* the measurements, assuming $\mathbf{Q}\mathbf{d}(k) = \mathbf{0}$, as implemented in proposition 1, the implication of which is described in the following proof:

Proof of proposition 1. From assumption 3, we observe that the test statistic in (4) has the property $\|\mathbf{P}_{\mathbf{Q}\mathbf{c}}^\perp \mathbf{Q}\mathbf{d}(k)\|^2 = \mathbf{0}$, $\exists \mathbf{Q}^\top \mathbf{Q} \in \{\mathbf{P}_d^\perp\}$. It is expected that the above property is best satisfied, by a candidate \mathbf{Q} , when the goodness-of-fit of the test statistic with respect to its corresponding distribution, assuming $\mathbf{Q}\mathbf{d}(k) = \mathbf{0}$,

$$\frac{\left(\|\mathbf{P}_{\mathbf{Q}\mathbf{c}}^\perp \mathbf{Q}\mathbf{y}(k)\|^2 - \mathbb{E}\left[\|\mathbf{P}_{\mathbf{Q}\mathbf{c}}^\perp \mathbf{Q}\mathbf{y}(k)\|^2 \mid \mathbf{Q}\mathbf{d}(k) = \mathbf{0}\right]\right)^2}{\text{Cov}\left[\|\mathbf{P}_{\mathbf{Q}\mathbf{c}}^\perp \mathbf{Q}\mathbf{y}(k)\|^2 \mid \mathbf{Q}\mathbf{d}(k) = \mathbf{0}\right]}$$

is minimized. We conclude the proof by observing that when $\mathbf{Q}\mathbf{d}(k) = \mathbf{0}$, it is true that $\|\mathbf{P}_{\mathbf{Q}\mathbf{c}}^\perp \mathbf{Q}\mathbf{y}(k)\|^2$ has a central chi-square distribution with central moments:

$$\begin{aligned}\mathbb{E}\left[\|\mathbf{P}_{\mathbf{Q}\mathbf{c}}^\perp \mathbf{Q}\mathbf{y}(k)\|^2 \mid \mathbf{Q}\mathbf{d}(k) = \mathbf{0}\right] &= \|\mathbf{Q}\|_0 - 1 \\ \text{Cov}\left[\|\mathbf{P}_{\mathbf{Q}\mathbf{c}}^\perp \mathbf{Q}\mathbf{y}(k)\|^2 \mid \mathbf{Q}\mathbf{d}(k) = \mathbf{0}\right] &= 2(\|\mathbf{Q}\|_0 - 1).\end{aligned}$$

□

State Estimation: Given the expected null space of the corruption, $\hat{\mathbf{P}}_d^\perp$, from proposition 1, the minimum mean-squared error estimate of the state is

$$\mathbb{E}\left[x(k) \mid \mathbf{P}_{\mathbf{Q}\mathbf{c}}^\perp \mathbf{Q}\mathbf{y}(k)\right] = \frac{\mathbf{c}^\top \mathbf{Q}^\top \mathbf{Q}\mathbf{y}(k)}{\mathbf{c}^\top \mathbf{Q}^\top \mathbf{Q}\mathbf{c}}$$

where $\mathbf{Q}^\top \mathbf{Q} = \hat{\mathbf{P}}_d^\perp$.

3.2 Finite-Horizon Parameter-Invariant Control

The parameter-invariant controller utilizes a time-history of the state estimates and control inputs to design a finite-horizon controller that stabilizes the mean and is invariant to the unknown system dynamics.

Proof of proposition 2. Defining $\hat{x}(k) = \mathbb{E}[x(k)|\mathbf{y}(k)]$, we consider the following time-series generated by (1),

$$\begin{aligned}
& \begin{bmatrix} \hat{x}(k_0) \\ \hat{x}(k_0 + 1) \\ \vdots \\ \hat{x}(k) \\ \alpha\hat{x}(k) \\ \vdots \\ \alpha^{\kappa-k}\hat{x}(\kappa) \end{bmatrix} = \begin{bmatrix} 1 & & & & & \\ -a & 1 & & & & \\ & \ddots & \ddots & & & \\ & & -a & 1 & & \\ & & & -a & 1 & \\ & & & & \ddots & \ddots \\ & & & & & -a & 1 \end{bmatrix}^{-1} \begin{bmatrix} \hat{x}(k_0) \\ \mu + bu(k_0) \\ \vdots \\ \mu + bu(k-1) \\ \mu + bu(k) \\ \vdots \\ \mu + bu(\kappa-1) \end{bmatrix} \\
\iff & \begin{bmatrix} \hat{x}(k_0) \\ \hat{x}(k_0 + 1) \\ \vdots \\ \hat{x}(k) \\ \alpha\hat{x}(k) \\ \vdots \\ \alpha^{\kappa-k}\hat{x}(\kappa) \end{bmatrix} = \begin{bmatrix} 1 & & & & -a & \\ -a & 1 & & & & \\ & \ddots & \ddots & & & \\ & & -a & 1 & & \\ & & & -a & 1 & \\ & & & & \ddots & \ddots \\ & & & & & -a & 1 \end{bmatrix}^{-1} \begin{bmatrix} \hat{x}(k_0) - a\alpha^{\kappa-k}\hat{x}(k) \\ \mu + bu(k_0) \\ \vdots \\ \mu + bu(k-1) \\ \mu + bu(k) \\ \vdots \\ \mu + bu(\kappa-1) \end{bmatrix} \\
\iff & \mathbf{s} = \bar{\mathbf{V}}^\top (\mathbf{I} + a\mathbf{A})^{-1} \mathbf{V} \begin{bmatrix} \hat{x}(k_0) - a\alpha^{\kappa-k}\hat{x}(k) \\ \mathbf{b}\mathbf{u}_{k_0:\kappa-1} + \mu\mathbf{1} \end{bmatrix} \\
\iff & \mathbf{V}\mathbf{s} = (\mathbf{I} + a\mathbf{A})^{-1} \mathbf{V} \begin{bmatrix} \hat{x}(k_0) - a\alpha^{\kappa-k}\hat{x}(k) \\ \mathbf{b}\mathbf{u}_{k_0:\kappa-1} + \mu\mathbf{1} \end{bmatrix} \\
\iff & \mathbf{0} = \mathbf{H}\boldsymbol{\theta} + \mathbf{V}_{1:\kappa-k_0} \mathbf{u}_{k_0:\kappa-1} \\
\iff & \mathbf{P}_H^\perp \mathbf{V}_{1:\kappa-k_0} \mathbf{u}_{k_0:\kappa-1} = \mathbf{0} \quad (\text{by Nyquist-Shannon Theorem})
\end{aligned}$$

where

$$\boldsymbol{\theta} = \left[-\frac{a}{b}, -\frac{1}{b}, \frac{\mu}{b}, \frac{x(k_0) - ax(\kappa) - \mu}{b} \right]^\top.$$

The first implication transforms the time-series signal into an equivalent $(\kappa - k_0 + 1)$ -step periodic signal by augmenting the initial condition. By the Nyquist-Shannon Theorem, the future control inputs, $\mathbf{u}_{k:\kappa-1}$, can be reconstructed from the previous measurements, a desired state trajectory, and the past control inputs, $\mathbf{u}_{k_0:k-1}$ if $\kappa - k_0 < 2(k - k_0) - \|\Theta\|_0$. Since this criteria is satisfied in the worst-case by assuming all parameters are non-zero, the control sequence that generates the desired future state trajectory is reconstructed by solving

$$\mathbf{P}_H^\perp \mathbf{V}_{1:\kappa-k_0} \mathbf{u}_{k_0:\kappa-1} = \mathbf{0}$$

assuming $\kappa - k_0 < 2(k - k_0) - 4$. \square

Recalling that the controller is design to stabilize the mean through an additive control input, the predicted covariance of the state is given in the following corollary:

Corollary 1 - Covariance of Predicted State: Assuming $\mathbf{Q}^\top \mathbf{Q} \in \{\hat{\mathbf{P}}_d^\perp\}$, then

$$\text{Cov} \left[x(t) | \mathbf{P}_{\mathbf{Q}c}^\perp \mathbf{Q} \mathbf{y}(k) \right] = \text{Cov} \left[x(k) | \mathbf{P}_{\mathbf{Q}c}^\perp \mathbf{Q} \mathbf{y}(k) \right] + a^{t-k-1} \sigma, \quad \forall t \geq k$$

Satisfying the equality constraint in proposition 2 restricts the horizon for which the controller can be designed. As a best estimate of the control sequence which satisfies proposition 2, we employ a maximum likelihood estimate of the future control sequence as

$$\mathbf{u}_{k:\kappa-1} = -\mathbf{G} \mathbf{u}_{k_0:k-1} \quad (5)$$

where

$$\mathbf{G} = \left(\bar{\mathbf{V}}_{k-k_0+1:\kappa}^\top \mathbf{P}_H^\perp \mathbf{V}_{k-k_0+1:\kappa} \right)^{-1} \bar{\mathbf{V}}_{k-k_0+1:\kappa}^\top \mathbf{P}_H^\perp \mathbf{V}_{1:k-k_0} \quad (6)$$

The resilient sensor fusion strategy presented in this section identifies a maximum likelihood estimate of the corruption null space, consistent with commonly adopted goodness-of-fit approaches. The resulting estimator is employed to design an parameter-invariant control sequence which stabilizes the mean of the estimate at a convergence rate of α .

4 Resilient Parameter-Invariant Controller Design

From the previous section, it is clear that the performance of the resilient state estimator affects the controller. Since it is a primary concern to secure the measurements against malicious attacks, and a secondary concern to maximize the performance with respect to the environmental unknowns, we introduce the following proposition to design the parameter-invariant controller convergence rate, α :

Proposition 3 - Controller Design : Given problem 1, $\hat{\mathbf{P}}_d^\perp$ from proposition 1, then a control sequence \mathbf{u} in proposition 2 designed assuming α will satisfy (2) if and only if

$$0 \leq \alpha \leq \sqrt{\frac{\alpha^* - \frac{\sigma + \frac{1}{\mathbf{c}^\top \hat{\mathbf{P}}_d^\perp \mathbf{c}}}{\eta}}{\frac{1}{\eta \mathbf{c}^\top \hat{\mathbf{P}}_d^\perp \mathbf{c}} + 1}}$$

Proof. We define $\mathbf{l} = \hat{\mathbf{P}}_d^\perp \mathbf{c} (\mathbf{c}^\top \hat{\mathbf{P}}_d^\perp \mathbf{c})^{-1}$, and write

$$\begin{aligned} \alpha \leq \sqrt{\frac{\alpha^* - \frac{\sigma + \mathbf{l}^\top \mathbf{l}}{\eta}}{\frac{1}{\eta} \mathbf{l}^\top \mathbf{l} + 1}} &\iff \alpha^2 \left(\frac{\mathbf{l}^\top \mathbf{l}}{\eta} + 1 \right) + \frac{\sigma + \mathbf{l}^\top \mathbf{l}}{\eta} \leq \alpha^* \\ \iff \alpha^2 \left(\mathbf{l}^\top \mathbf{l} + \|x(k)\|^2 \right) + \left(\mathbf{l}^\top \mathbf{l} + \sigma \right) &\leq \alpha^* \|x(k)\|^2, \quad \forall \|x(k)\|^2 \geq \eta \\ \iff \mathbb{E} \left[\|\alpha \mathbb{E} [x(k) | \mathbf{y}(k)]\|^2 + \text{Cov} [x(k+1) | \mathbf{y}(k)] \right] &\leq \alpha^* \|x(k)\|^2, \quad \forall \|x(k)\|^2 \geq \eta \\ \iff \mathbb{E} \left[\|\mathbb{E} [x(k+1) | \mathbf{y}(k)]\|^2 + \text{Cov} [x(k+1) | \mathbf{y}(k)] \right] &\leq \alpha^* \|x(k)\|^2, \quad \forall \|x(k)\|^2 \geq \eta \\ \iff \mathbb{E} [\|x(k+1)\|^2] \leq \alpha^* \|x(k)\|^2, \quad \forall \|x(k)\|^2 \geq \eta \end{aligned}$$

□

A direct consequence of proposition (3) is the following corollary identifying when a resilient parameter-invariant controller exists that satisfies the performance criteria in (2).

Corollary 2 - Controller Existence : A resilient parameter-invariant controller satisfying (2) exists if and only if

$$\frac{1}{\mathbf{c}^\top \hat{\mathbf{P}}_d^\perp \mathbf{c}} \leq \alpha^* \eta - \sigma$$

By applying Markov's inequality to proposition 3, a probabilistic bound on the likelihood the state diverges is provided in the following corollary:

Corollary 3 - Probability of Divergence: Assuming proposition 3,

$$\Pr [\|x(k+1)\|^2 \geq \|x(k)\|^2] \leq \alpha^*, \quad \forall \|x(k)\|^2 \geq \eta \quad (\text{Markov's inequality})$$

The resilient parameter-invariant controller formulated in this section is evaluated through simulation in the following section.

5 Simulation Results

This section provides a qualitative evaluation of the resilient parameter-invariant controller. This evaluation is presented in two subsections. The following subsection presents disturbance rejection results considering first order systems, both unstable and stable, when sensor corruption is both present and absent. The final subsection presents a simulated cruise control scenario for robotic platforms.

5.1 Disturbance Rejection

In this subsection, we evaluate the resilient parameter-invariant controller with respect to disturbance rejection. We consider the following global variables for all simulations in this subsection

- $b = 1$ (input gain);
- $C = [2, 3, 4, 5, 10]^T$ (measurement gain);
- $x(0) = 5$ (initial condition);
- $\mu = -1$ (process noise mean);
- $\sigma = 0.01$ (process noise covariance);
- $\alpha^* = 0.1, \eta = 1$ (stochastic boundedness parameters);

Additionally, we use a windowed approach for the controller design, where we use the last 20 measurements (19 previous inputs), to design a 10-step finite horizon control sequence at each time step. Assuming this system and controller design strategy, we evaluate the resilient controller when the system is stable ($a = 0.8$) and unstable ($a = 1.8$) when all the sensors are uncorrupted, and when the first and last sensors are corrupted by a random attacker using a zero-mean Gaussian attack with variance 10. The remainder of this subsection addresses these four scenarios: (a) stable, no corruption; (b) stable, with corruption; (c) unstable, no corruption; (d) unstable, with corruption.

All figures in this section (i.e. Fig. 1 - 4), have the same subplot structure, as demonstrated by Fig. 1. The upper most subplot denotes the state trajectory (i.e. the value of $x(k)$ during the entire simulation). The second subplot denotes the two-norm of the state trajectory, $\|x\|$, as a solid line, the crosses at each time step identify the two-norm of the state-estimate, and the dashed line denotes the stochastic boundedness threshold, η . The third subplot identifies the control input, u , at each time step. The fourth subplot identifies whether a sensor (1 to 5) was selected at a specific time step (denoted by a cross) or identified as corrupted at each time step (no cross). The final subplot illustrates the convergence rate, α , chosen by the parameter-invariant controller at each time step such that the desired convergence rate α^* is achieved.

Fig. 1 presents the results for a stable system with no sensor attacks. In Fig. 1, we simulate a state disturbance injection (beyond the process noise) occurring every 30 time steps. It is clear that the controller is capable of stabilizing the system, and achieves the performance bound on the norm of the state; however, we note that despite no attacks being present in this simulation it is assumed that some of the sensor measurements are corrupted at most time steps. This results from the fact that unlike classical state-estimators, which minimize the mean-squared error of the state estimate, the resilient sensor fusion minimizes a normalized deviation of the measured sensor noise from a weighted average of the sensor measurements. This results in the rejection of measurements which are significantly different from the mean of the sensor measurements. As a consequence, depending on the specific subset of measurements accepted as secure, the controller convergence rate changes to maintain the performance specification.




Fig. 1. Stable system ($a = 0.8$) with no sensor attacks.

Simulated results for a stable system in the presence of sensor attacks is presented in Fig. 2. We consider the same state disturbance injection as in Fig. 1, and observe a very similar performance in terms of disturbance rejection. However, this comes at the cost of an increased variance in the controller sequence (as compared to Fig. 1). The increased variance is undesirable in most physical actuators since this results in increased strain and fatigue on mechanical parts. The increased variance is a direct result of designing the controller sequence to match the performance specification, and can be reduced by upper bounding the desired state convergence rate, α , employed by the parameter-invariant controller. Through a comparison between Fig. 1 and Fig. 2 we notice that sensor one is selected significantly less when under attack and sensor five is selected only marginally less, despite being corrupted by the same attack. This is expected since given the same attack, signal-to-corruption ratio is greater in sensor five than in sensor one. Recalling that the measurement model in (1) is normalized such that all the sensors have the same noise profile, sensor five having a higher




Fig. 2. Stable system ($a = 0.8$) with sensor 1 and 5 under attack.

signal-to-corruption ratio (assume the same attack) than sensor one is equivalent to stating that since sensor five is less noisy than sensor one, sensor five is more likely to be trusted.

A unstable system with no sensor attacks is considered in Fig. 3. Similar to the results for the stable system, here we observe that resilient parameter-invariant controller stabilizes the system with respect to unknown disturbances, and has a response very similar to the stable system results in Fig. 1. Although the control input is different, the performance is nearly identical, which is a direct result of the controller design requirement to be invariant to the unknown system parameters, regardless of stability. Similarly, we notice a similar sensor selection and rejection profile as in Fig. 1. This result is consistent with the fact that the sensor fusion strategy is design invariant to the unknown state, which orthogonalizes the issues of stability and resilience for the purposes of identifying corrupted sensors.




Fig. 3. Unstable system ($a = 1.8$) with no sensor attacks.

The final figure in this subsection, Fig. 4, illustrates the results when considering an unstable system and with sensor attacks. Consistent with the results when no sensor attacks are present, we observe nearly identical results when sensors are attacked, regardless of the underlying system stability. This further illustrates that the specification based-design of the resilient parameter-invariant controller can be achieved regardless of the underlying system parameters. We note, however, a multi-step design specification will vary with the system stability since the multi-step predicted state covariance is a non-linear function of the process noise and state dynamic gain, a . For this reason, it is necessary that the resilient parameter invariant controller be designed at each time step in order to satisfy the performance specification invariant to the unknown state dynamics.




Fig. 4. Unstable system ($a = 1.8$) with sensor 1 and 5 under attack .

5.2 Resilient Cruise Control

To evaluate the resilient parameter-invariant controller as a potential cruise control mechanism, we consider a simplified first-order model of a mobile robot as

$$\begin{aligned}
 x(k + 1) &= 0.95x(k) + u(k) + w(k) \\
 y(k) &= [2, 3, 4]^T x(k) + v(k)
 \end{aligned}
 \tag{7}$$

where $w(k) \in [-1, .01]$, and we assume $\alpha^* = .1$, and $\eta = 1$ for the performance constraint. The parameter choices in this simulation are motivated by the fact that the velocity of surface robots is accurately modeled as a stable first-order system when set-point control is performed at a rate an order of magnitude greater than the electro-mechanical system time-constant.

In this simulation, we assume the initial vehicle speed is 5 km/h ($x(0) = 5$) and at time step 50 the speed is desired to increase to 15 km/h. Additionally, we assume the control input is bounded on the interval of -1 to 5. The result




Fig. 5. Landshark cruise control with no corruption.

of employing the resilient parameter-invariant controller for cruise control is provided in Fig. 5 without sensor attacks. We observe in Fig. 5 that the vehicle speed is stabilized with a lag in the state trajectory. The lag is a result of the fact that a history of measurements is required to generate a control sequence. It remains a focus of future research to reduce this lag. A portion of the state trajectory lag is due to the saturation of the control signal, which does not affect the stability of the system despite the resilient control law being designed without a constraint on the control sequence.

The cruise controller performance in the presence of a sensor attack on sensor three is illustrated in Fig. 6. Here we observe, and consistent with previous observations, that when attacked, the measurement of sensor three is accepted less often. In this case, sensor one (known to be more noisy) is selected more often. This results in the need for, on a average, a more aggressive control law to meet the performance criteria.




Fig. 6. Landshark cruise control with corruption.

6 Discussion and Future Work

This work addresses the problem of resilient control of unknown first-order stochastic LTI systems in the presence of sensor attacks. A resilient sensor fusion strategy is introduced that minimizes the likelihood of a corrupted sensor being trusted. A time-series concatenation of secure estimates is utilized in cooperation with a desired state trajectory and previous control inputs to design a mean-stabilizing finite-horizon control sequence, invariant to the unknown system parameters. The parameter-invariant controller is designed using a frequency domain representation of an equivalent time-series representation of the system inputs to system outputs, thus allowing the design of a mean-stabilizing controller. Simulated results illustrate that the resilient parameter-invariant controller is capable of stabilizing unknown state disturbances and can perform state trajectory tracking.

Future work on this topic includes the extension of the scalar results in this paper to multi-dimensional systems with known dynamical structures but un-

known parameters. A limiting factor in the application of the resilient parameter-invariant controller is the requirement (at this time) that the system-to-be-controlled behave as a linear first order system. While most stable systems can (at some time scale) be approximated by a first order system, this approximation can be very poor. As future work, we plan to evaluate the proposed control strategy on various experimental platforms to determine the robustness of the parameter-invariant control with respect to the design dependency on a first-order linear system. Additionally, further insight is needed to investigate a method to reduce (or remove) the lag and to quantify its behavior in terms of the underlying system and design parameters. Experimentation of the resilient parameter-invariant controller is planned on a robotic platform as a potentially cooperative approach with model-based approaches which assume knowledge of model parameters.

7 Acknowledgements

This material is based on research sponsored by DARPA under agreement number FA8750-12-2-0247. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

References

1. Willsky, A.: A survey of design methods for failure detection in dynamic systems. *Automatica* **12** (1976) 601–611
2. Qiu, L.: Essentials of robust control Kemin Zhou, John C. Doyle Prentice-Hall, Englewood Cliffs, NJ, 1998, ISBN: 0-13-790874-1. *Automatica* **38**(5) (May 2002) 910–912
3. Astrom, K.J., Wittenmark, B.: Adaptive Control. 2nd edn. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA (1994)
4. Trees, H.L.V.: Detection, Estimation, and Modulation Theory. John Wiley & Sons, Inc., New York (1968)
5. Scharf, L.L.: Statistical Signal Processing, Detection, Estimation, and Time Series Analysis. Addison-Wesley Publishing Company Inc., Reading, Massachusetts (1991)
6. Schenato, L., Sinopoli, B., Franceschetti, M., Poolla, K., Sastry, S.: Foundations of control and estimation over lossy networks. *Proceedings of the IEEE* **95**(1) (jan. 2007) 163–187
7. Gupta, A., Langbort, C., Basar, T.: Optimal control in the presence of an intelligent jammer with limited actions. In: Decision and Control (CDC), 2010 49th IEEE Conference on. (dec. 2010) 1096–1101
8. Pasqualetti, F., Dörfler, F., Bullo, F.: Attack detection and identification in cyber-physical systems – part ii: Centralized and distributed monitor design. Technical Report arXiv:1202.6049 (Feb 2012)

9. Sundaram, S., Pajic, M., Hadjicostis, C.N., Mangharam, R., Pappas, G.J.: The wireless control network: Monitoring for malicious behavior. In: CDC. (2010) 5979–5984
10. Weimer, J., Kar, S., Johansson, K.H.: Distributed detection and isolation of topology attacks in power networks. In: Proceedings of the 1st international conference on High Confidence Networked Systems. HiCoNS '12, New York, NY, USA, ACM (2012) 65–72
11. Weimer, J., Ahmadi, S.A., Araujo, J., Mele, F.M., Papale, D., Shames, I., Sandberg, H., Johansson, K.H.: Active actuator fault detection and diagnostics in hvac systems. In: 4th ACM Workshop On Embedded Systems For Energy-Efficiency In Buildings (BuildSys), Toronto, Canada (2012)
12. Fawzi, H., Tabuada, P., Diggavi, S.N.: Secure estimation and control for cyber-physical systems under adversarial attacks. CoRR **abs/1205.5073** (2012)
13. Pintelon, R., Guillaume, P., Rolain, Y., Schoukens, J., Van hamme, H.: Parametric identification of transfer functions in the frequency domain—a survey. Automatic Control, IEEE Transactions on **39**(11) (nov 1994) 2245–2260
14. Ljung, L., ed.: System identification (2nd ed.): theory for the user. Prentice Hall PTR, Upper Saddle River, NJ, USA (1999)
15. Van Den Hof, P.: Identification and control - closed-loop issues. Automatica **31**(12) (December 1995) 1751–1770

Appendix: Notation and Definitions

This appendix introduces notation and definitions that prove useful in this work.

7.1 Notation

In this subsection, we illustrate the various variable notations using varying fonts and capitalization of the letter z :

- plain upper case italic fonts \rightarrow constant, Z ;
- plain lower case italic fonts \rightarrow scalar (or function with scalar range), z ;
- bold lower case italic fonts \rightarrow vector (or function with vectoral range), \mathbf{z} ;
- bold lower case plain fonts \rightarrow vector of concatenated vectors, \mathbf{z} ;
- bold upper case italic fonts \rightarrow matrix, \mathbf{Z} ;

For vectors we write \mathbf{z}_i to denote the i -th position of \mathbf{z} and $\mathbf{z}_{i:j}$ to be the sub-vector of \mathbf{z} consisting of the i -th through j -th elements, inclusively. Similarly, for vectors of vectors we write \mathbf{z}_i to denote the i -th sub-vector and $\mathbf{z}_{i:j}$ to be the sub-vector of \mathbf{z} consisting of the i -th through j -th sub-vectors, inclusively. Lastly, for matrices we write \mathbf{Z}_i to be the i -th column of \mathbf{Z} and $\mathbf{Z}_{i:j}$ to be the sub-matrix consisting of the i -th through j -th columns, inclusively.

We use the notation $\Pr[x|y]$ and $\mathbb{E}[x|y]$ to denote the probability of x given y and the expected value of x given y , respectively,

7.2 Definitions

This subsection defines matrices and constants which appear in the remainder of this work. Consistent with the previous subsection's use of the letter z to illustrate various properties, we write \mathbf{P}_Z to be the projection matrix corresponding to the general matrix \mathbf{Z} and \mathbf{P}_Z^\perp to be the projection matrix corresponding to the null-space of \mathbf{Z} ,

$$\mathbf{P}_Z = \mathbf{Z} \left(\bar{\mathbf{Z}}^\top \mathbf{Z} \right)^{-1} \bar{\mathbf{Z}}^\top \quad \text{and} \quad \mathbf{P}_Z^\perp = \mathbf{I} - \mathbf{P}_Z,$$

where $\bar{\mathbf{Z}}^\top$ is notation denoting the transpose of the complex conjugate of \mathbf{Z} and \mathbf{I} the identity matrix.

For an arbitrary positive integer $N \in \mathbb{N}^+$, $\omega = \frac{2\pi\sqrt{-1}}{N}$, $\lambda(n) = \mathbf{e}^{\omega n}$, we define the matrices $\mathbf{V}, \mathbf{A} \in \mathbb{C}^{N \times N}$ as

$$\begin{aligned} \mathbf{A} &= \text{diag}[\lambda(0), \dots, \lambda(N-1)] \\ \mathbf{V}_n &= \frac{1}{\sqrt{N}} \left[1, (\lambda(n))^1, \dots, (\lambda(n))^{N-1} \right]^\top \end{aligned}$$

and note that \mathbf{V} is the normalized N -point Discrete Fourier Transform (DFT) matrix, such that

$$\bar{\mathbf{V}}^\top \mathbf{V} = \mathbf{V} \bar{\mathbf{V}}^\top = \mathbf{I}$$