



Research article

PLATE: A perception-latency aware estimator[☆]Rodrigo Aldana-López, Rosario Aragüés, Carlos Sagüés^{*}*Departamento de Informática e Ingeniería de Sistemas (DIIS) and Instituto de Investigación en Ingeniería de Aragón (I3A), Universidad de Zaragoza, Zaragoza 50018, Spain*

ARTICLE INFO

Article history:

Received 20 July 2022

Received in revised form 25 July 2023

Accepted 11 August 2023

Available online 14 August 2023

Keywords:

Target-tracking

Perception-latency

State estimation

ABSTRACT

Target tracking is a popular problem with many potential applications. There has been a lot of effort on improving the quality of the detection of targets using cameras through different techniques. In general, with higher computational effort applied, i.e., a longer perception-latency, a better detection accuracy is obtained. However, it is not always useful to apply the longest perception-latency allowed, particularly when the environment does not require to and when the computational resources are shared between other tasks. In this work, we propose a new Perception-LATency aware Estimator (PLATE), which uses different perception configurations in different moments of time in order to optimize a certain performance measure. This measure takes into account a perception-latency and accuracy trade-off aiming for a good compromise between quality and resource usage. Compared to other heuristic frame-skipping techniques, PLATE comes with a formal complexity and optimality analysis. The advantages of PLATE are verified by several experiments including an evaluation over a standard benchmark with real data and using state of the art deep learning object detection methods for the perception stage.

© 2023 The Authors. Published by Elsevier Ltd on behalf of ISA. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

The Target Tracking (TT) problem involves using vision or range sensors to locate and track the target's trajectory. Sequential or online tracking aims to provide real-time estimates of the target's state as measurements are obtained [1]. This problem has a variety of applications, including automated surveillance [2], crowd analysis and pedestrian intention prediction [3], gesture recognition [4], and traffic monitoring [5]. TT is also important in the robotics community, where it is used in robot trajectory tracking [6], cooperative surveillance [7], and object tracking for localization and mapping [8]. Most TT solutions rely on a detector, which can be based on traditional vision methods [9] or Deep Learning (DL) techniques [10,11]. When multiple targets are involved, an association and re-identification stage is added to distinguish detections between targets. Furthermore,

probabilistic inference techniques such as Kalman filtering, particle filtering, or other Bayesian methods are typically used for model-based tracking [1].

In this work, we adopt the term perception-latency from [12] to refer to the time elapsed from the instant the sensors start sampling to the moment the target is detected and identified from the raw sensor data. Note that the perception-latency can be varied by changing the perception configuration, which includes the resolution of visual or range sensors, the number of features considered in the perception algorithm, and model hyperparameters if DL detectors are used. In many TT approaches, the perception configuration can be modified to improve accuracy at the expense of a longer perception-latency [13], leading to a perception-latency/accuracy trade-off. For radar-based perception, the term dwell-time refers to the time taken by the radar to acquire a measurement [14], which is considered part of the perception latency and introduces a similar trade-off between quality and duration.

Other works, such as [15], have studied a similar latency/accuracy trade-off in networked control systems where the latency arises from the communication channel. While there are some similarities between this problem and the present work, our interest lies in the setting where the sensor, perception algorithm, and estimation algorithm are all in the same computing unit. In this case, the latency due to information transmission is negligible when compared to the perception latency, which is the primary concern.

[☆] This work was supported by projects PID2021-124137OB-I00 and TED2021-130224B-I00 (MCIN/AEI/10.13039/501100011033, by ERDF A way of making Europe and by the European Union NextGenerationEU/PRTR), DGA T45_23R (Gobierno de Aragón, Spain), 739841 (Consejo Nacional de Ciencia y Tecnología CONACYT-Mexico), and by Universidad de Zaragoza, Banco Santander.

^{*} Correspondence to: Universidad de Zaragoza: Escuela de Ingeniería y Arquitectura, C. María de Luna, 3, Zaragoza 50018, Spain.

E-mail addresses: rodrigo.aldana.lopez@gmail.com, 794214@unizar.es (R. Aldana-López), raragues@unizar.es (R. Aragüés), csagues@unizar.es (C. Sagüés).

In addition, real-time tracking is often required, meaning that the perception-latency is subject to a real-time constraint determined by the sampling period of the vision sensors, typically 1/30 s or 1/60 s for traditional vision cameras [16]. Real-time TT is particularly relevant for resource-constrained platforms such as mobile systems [17] and smart cameras [18]. The perception-latency/accuracy trade-off is typically addressed by selecting the optimal perception configuration that satisfies the real-time constraint. Most approaches using DL involve testing different neural network architectures to find a suitable accuracy compromise with smaller models and less perception-latency.

Moreover, note that perception is not always the only task running in the system. Robotic platforms use computing resources for motion planning, control and coordination tasks. Hence, setting the real-time constraint for the perception task as the full sampling period of the vision sensor is still unrealistic. In addition, as pointed out in [19], it is not always necessary to run the perception task for each new frame if the dynamics of the scene does not require to. Instead, one may use a frame-skipping policy obtaining a better resource usage. This means that CPU load for vision-based perception must be managed as well.

In this sense, looking for adaptive perception schemes, in which perception-latency changes according to the situation, might be beneficial for a better resource usage under the perception-latency/accuracy trade-off. Some DL techniques allow adaptive forward latency such as the so called Anytime Neural Networks (ANN). ANNs provide a way to select online between different compromises of latency and accuracy [20]. These type of networks were used in [21] in order to schedule different DL jobs with latency chosen such that real-time deadlines are maintained with the best possible accuracy overall. However, being a general DL scheduler, [21] does not take into account that for a tracking problem, while a longer latency might lead to a better accuracy for the detector, the model uncertainty might increase as well due to the delay introduced by the perception-latency [22].

Instead of using an adaptive perception-latency, some works have used frame-skipping techniques in order to reduce computational load. In [19], the authors propose an event-triggered rule in which the detection is only updated under certain events, based on the expected motion of the target. This approach reduces the computational burden of the perception task by effectively skipping frames when the expected motion of the target does not require to. Other approaches follow a similar idea such as the one in [18] in which the perception task goes to an idle state depending on the expected state of the environment, obtaining an energy efficient use of battery powered smart cameras. Similarly, in [23] a neural network is trained based on examples in order to make a decision on whether or not the detection for some target should be updated. Although some of these methods shows good experimental performance in certain scenarios, they are mostly based in several heuristic rules from which it is hard to provide formal guarantees on the impact over the resource usage and quality for TT in the general case.

In addition, the related problem of multi-sensor scheduling problem can be relevant in the perception scheduling context [24, 25]. In this problem, multiple sensors are available to measure the state of a dynamical system and the goal is to schedule them for control or estimation purposes under a given cost function. As discussed in [26] this problem is usually subject to a combinatorial explosion in terms of the number of available sensors and the time horizon for the cost function. Hence, most efforts in the literature have focused on obtaining efficient sub-optimal solutions. For instance, [27] take advantage of the structure of the cost function and the system to prune many of the possible schedules. Despite this, while this approach reduces the complexity, the problem may still be intractable for big horizon sizes.

Other approaches adopt a greedy policy which picks the best sensor at each time which optimizes a one-step look-ahead cost as in [28]. Nonetheless, despite the simplicity of this approach, the result tends to be quite conservative since the behavior on the long run is not taken into account.

On the other hand, another approach is to avoid the combinatorial explosion by restricting only to periodic schedules as in [26, 29, 30]. However, this is conservative as well, being sub-optimal when not in steady state. Moreover, in most cases the actual optimal period for the schedule is unknown and depends on the particular problem setting. Other efficient approaches were proposed in [31, 32] which do not require the periodic schedule assumption. Nonetheless, the actual form of the scheduling policies in the previous works is tightly coupled to the particular structure of the systems under consideration and the cost function. In this sense, it is not possible to apply directly these ideas when we take into account the perception-latency trade-off as well as computational load and energy.

Motivated by this, our contribution is a new Perception-LATency aware Estimator (PLATE). PLATE schedules different perception configurations in different moments of time in order to optimize a performance measure that takes into account the perception-latency/accuracy trade-off aiming to improve estimation quality and resource usage. We show that the proposed perception-latency scheduling policy is not subject to a combinatorial explosion and approximates the optimal perception schedule with arbitrary precision as more computing time is employed. The advantages of PLATE over the state of the art are summarized as follows

- Compared to the rest of the literature, PLATE explicitly takes into account the effect of the delay introduced by the perception-latency on the target uncertainty model. Moreover, the algorithm is not tightly coupled to the form of the cost function, making it more general and versatile than related approaches. This allows the use of PLATE in an online setting even when the perception quality is estimated to be different from its nominal values.
- Compared to existing heuristic frame-skipping techniques [1, 18, 19], we provide a theoretical analysis of our proposal, equipped with formal guarantees under linear model and additive Gaussian disturbance assumptions.
- Compared to multi-sensor scheduling approaches, we do not assume periodicity of the schedule, allowing to less conservative results, which can adapt to varying perception quality.

The advantages of PLATE are verified by several experiments including an evaluation over a standard benchmark with real data and using state of the art DL methods for the perception stage.

1.1. Notation

Let $\bar{\mathbb{R}}_{\geq 0} = \mathbb{R}_{\geq 0} \cup \{\infty\} \geq 0 \geq 0$. Moreover, let $\mathbb{E}\{\bullet\}$ be the expectation operator and $\text{cov}\{\bullet, \bullet\} = \mathbb{E}\{(\bullet - \mathbb{E}\{\bullet\})(\bullet - \mathbb{E}\{\bullet\})^T\}$ and $\text{cov}\{\bullet\} := \text{cov}\{\bullet, \bullet\}$. Let $\text{tr}(\bullet)$ and $\text{vec}(\bullet)$ be the trace and vectorization operators respectively. Let \otimes denote the Kronecker product. Let $\text{diag}(v_1, \dots, v_n) \in \mathbb{R}^{n \times n}$ be a matrix with diagonal components $v_1, \dots, v_n \in \mathbb{R}$. We use $\mathbf{x}(t)$ to denote the evaluation of a signal \mathbf{x} at continuous time t whereas $\mathbf{x}[k] := \mathbf{x}(\tau_k)$ for discrete time instants $\{\tau_k\}_{k=0}^{\infty}$. In addition, for a sequence $p = \{p_k\}_{k=0}^{\ell-1}$ of length $\ell \in \mathbb{N} \cup \{\infty\}$, denote $\text{len}(p) := \ell$. For any matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with components $[\mathbf{A}]_{ij} \in \mathbb{R}$, let $\|\mathbf{A}\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^n [\mathbf{A}]_{ij}^2}$ its Frobenius norm [33, Page 341]. Let $\mathbf{A} \succ \mathbf{B}$ or $\mathbf{A} \succeq \mathbf{B}$ denote positive definiteness and semi-definiteness of $\mathbf{A} - \mathbf{B}$ respectively. For a set \bullet , let $\min\{\bullet\}$, $\max\{\bullet\}$, $\inf(\bullet)$, $\sup(\bullet)$ represent the minimum, maximum, infimum and supremum standard set operators.

2. Problem statement

Consider the following target model, with state $\mathbf{x}(t) \in \mathbb{R}^{n_x}$ e.g. containing its position and velocity. Assume a simplified motion model given by the following Stochastic Differential Equation (SDE):

$$d\mathbf{x}(t) = \mathbf{A}\mathbf{x}(t)dt + \mathbf{B}d\mathbf{w}(t), \quad t \geq 0 \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{n_x \times n_x}$, $\mathbf{B} \in \mathbb{R}^{n_x \times n_w}$ and $\mathbf{w}(t)$ is a n_w -dimensional Wiener processes with covariance $\text{cov}\{\mathbf{w}(s), \mathbf{w}(r)\} = \mathbf{W} \min(s, r)$ [34, Page 63]. As usual, the process $\mathbf{w}(t)$ models disturbances, unknown inputs for the target and non-modeled dynamics. Moreover, $\mathbf{x}(0)$ is normally distributed with mean \mathbf{x}_0 and covariance \mathbf{P}_0 .

The goal is to construct an estimation framework for the state $\mathbf{x}(t)$ using available sensors, e.g. vision or range. The sensors can produce *raw measurements* with a minimum sampling period of Δ_s . In order to use these measurements, they must be processed. For each raw measurement, the system can choose between D different perception methods which take the raw measurements to produce a *processed measurement* for the position of the target through a detection process. Thus, processed measurements are available to be used by the system only at, perhaps non-uniform, instants $\{\tau_k\}_{k=0}^\infty$ all integer multiples of Δ_s with $\tau_0 = 0$. Processed measurements are represented by $\mathbf{C}\mathbf{x}[k]$ with some constant matrix $\mathbf{C} \in \mathbb{R}^{n_z \times n_x}$ with (\mathbf{A}, \mathbf{C}) observable. Each method has a different perception-latency in $\{\Delta^1, \dots, \Delta^D\}$. These latencies are multiples of the sampling period Δ_s of the sensors leading to a frame-skipping technique. Hence, if method $p_k \in \{1, \dots, D\}$ is chosen at $t = \tau_k$, a new measurement $\mathbf{z}[k] = \mathbf{C}\mathbf{x}[k] + \mathbf{v}[k]$ is available at $t = \tau_k + \Delta^{p_k}$ where $\mathbf{v}[k]$ is a noise for the accuracy of the perception method.

The perception-latency/accuracy for the perception method is modeled by the covariance matrix $\mathbf{R}^{p_k} = \text{cov}\{\mathbf{v}[k]\}$. Typically, increasing the computing time Δ^{p_k} leads to a decrease in \mathbf{R}^{p_k} and an improvement in estimation precision. However, we do not make any assumptions about the specific perception-latency and quality model. Instead, for the sake of flexibility, a nominal \mathbf{R}^{p_k} is expected to be estimated based on the performance of the perception methods for the particular application of interest.

The main issue we address in this work is to obtain an estimate of the target state $\mathbf{x}(t)$ at any time $t \geq 0$ and choose which perception method to use at each time slot $[\tau_k, \tau_{k+1})$ provided some performance measure is optimized. In order to provide a concrete definition of the performance measure used in this work, consider the following definitions.

Definition 1 (Perception Schedules). A perception schedule p is a sequence of perception methods, i.e. $p = \{p_k\}_{k=0}^{\text{len}(p)-1}$ with $p_k \in \{1, \dots, D\}$ for some $\text{len}(p) \in \mathbb{N} \cup \{\infty\}$.

Definition 2 (Attention of a Perception Schedule). Let $\mathcal{I} \subset \mathbb{R}_+$ be an interval with $\sup(\mathcal{I}) - \inf(\mathcal{I}) > \max\{\Delta^1, \dots, \Delta^D\}$. Thus, the attention $\text{att}(p; \mathcal{I})$ of a perception schedule p for \mathcal{I} corresponds to the amount of processed measurements generated by p in such interval.

Thus, consider the following problem:

Problem 1 (Perception Scheduling). Design a causal estimator which picks a perception method p_k at each $t = \tau_k$ and produces an estimate $\hat{\mathbf{x}}(t)$ of $\mathbf{x}(t)$ with $\hat{\mathbf{P}}(t) := \text{cov}\{\mathbf{x}(t) - \hat{\mathbf{x}}(t)\}$ using only information prior to the instant t . Moreover, the estimator output must minimize

$$\mathcal{J}(\hat{\mathbf{x}}, p) = \frac{1}{T_f} \left(\int_0^{T_f} \text{tr}(\hat{\mathbf{P}}(t))dt \right) + \frac{\lambda_\alpha}{T_f} \left(\sum_{k=0}^{\alpha} r^{p_k} \right) \quad (2)$$

for some window of interest $[0, T_f]$, where $\alpha := \text{att}(p; [0, T_f])$, $\lambda_\alpha > 0$ and $r^{p_k} \geq 0$ are additional penalties assigned to each perception method.

Remark 1. The cost function in (2) is used to capture the perception latency-accuracy trade-off as explained in the following: the integral term measures the mean expected squared error of $\hat{\mathbf{x}}(t)$ by evaluating $\mathbb{E}(\hat{\mathbf{x}}(t) - \mathbf{x}(t))^T(\hat{\mathbf{x}}(t) - \mathbf{x}(t)) \equiv \text{tr}(\hat{\mathbf{P}}(t))$. Thus, the integral term in (2) represents the quality of the estimation $\hat{\mathbf{x}}(t)$. The summation term in (2) does not depend on the estimations $\hat{\mathbf{x}}(t)$ but only on the penalties r^{p_k} . If $r^{p_k} = 1$ for all $k = 0, 1, \dots$, this term penalizes the attention of p . In some scenarios, it is beneficial to keep this quantity small due to energy consumption. Similarly, non-skipped frames in vision-based perception require an exchange of information between a sensor and the computing unit, and a sensor sampling request. Therefore, processing fewer measurements is desirable for energy efficiency and to minimize the use of I/O buses in the system.

Remark 2. The latency Δ^{p_k} is usually not exclusively used to process sensor information, but the method also frees the computing unit for an interval of length $(1 - f^{p_k})\Delta^{p_k}$ with $f^{p_k} \in (0, 1)$. Thus, the CPU load in the interval $[0, T_f]$ can be expressed as $(1/T_f) \sum_{k=0}^{\alpha} f^{p_k} \Delta^{p_k}$. Therefore, penalties $r^{p_k} = f^{p_k} \Delta^{p_k}$ in (2) penalize CPU load.

The concepts of attention and CPU load, as discussed in the context of vision-based perception, share similarities with revisit and dwell-times in radar-based perception. Revisit time in a radar corresponds to the interval between consecutive detections, while dwell-time represents the time it takes for the sensor to generate a detection [14]. Both times have an impact on detection accuracy and energy consumption. Thus, revisit and dwell-times play similar roles to the inverse of attention and CPU load.

We propose PLATE as a strategy to compute the optimal schedule p and the optimal estimations $\hat{\mathbf{x}}(t)$ in the sense of Problem 1. Causality in Problem 1 is required in a sequential TT context. Thus, PLATE works as a causal estimator as is summarized in Algorithm 1 and is comprised by the combination of an estimation stage and a perception-latency scheduling policy.

Algorithm 1: PLATE loop

```

1  $\tau_s \leftarrow 0, k \leftarrow 0$ 
2 At sampling event  $t = \tau_s$  do
3   Estimation stage: Use  $\mathbf{x}_0$  and  $\{\mathbf{z}[0], \dots, \mathbf{z}[k-1]\}$  to
     compute  $\hat{\mathbf{x}}[k]$ .
4   Read raw data from sensors.
5   Perception latency scheduling: Decide which
     perception method  $p_k \in \{1, \dots, D\}$  to use.
6   Wait  $\Delta^{p_k}$  units of time until the perception-latency
     has elapsed to produce  $\mathbf{z}[k]$ , i.e.  $\tau_s \leftarrow \tau_s + \Delta^{p_k}$ .
7    $k \leftarrow k + 1$ 
8 end
```

3. Perception-latency aware estimation

In this section, we establish the structure of the estimation stage of PLATE. To do so, and in order to study the latency-precision trade-off under a cost of the form (2), it is useful to study an equivalent model of (1) as a sampled-data system. The following result follows from [35, Section 4.5.2]:

Proposition 1. Consider a perception schedule p . Hence, the solution $\mathbf{x}(t)$ of (1) satisfy:

$$\mathbf{x}(t) = \mathbf{A}_d(t - \tau_k)\mathbf{x}[k] + \mathbf{w}_d(t) \quad (3a)$$

$$\mathbf{x}[k+1] = \mathbf{A}_d(\Delta^{p_k})\mathbf{x}[k] + \mathbf{w}_d[k] \quad (3b)$$

for $t \in [\tau_k, \tau_{k+1})$, $k \geq 0$ where $\mathbf{A}_d(t - \tau_k) := \exp(\mathbf{A}(t - \tau_k))$ and $\mathbf{w}_d(t)$ normally distributed with $\text{cov}\{\mathbf{w}_d(t)\}$ given as

$$\mathbf{W}_d(t - \tau_k) = \int_0^{t-\tau_k} \mathbf{A}_d(\tau) \mathbf{B} \mathbf{W} \mathbf{B}^\top \mathbf{A}_d(\tau)^\top d\tau \quad (4)$$

The discrete-time nature of the perception process is made explicit by (3), which is equivalent to (1). As a result, we adopt (3) as the motion model throughout the manuscript.

The following result shows the computation of an optimal estimation $\hat{\mathbf{x}}(t)$ of $\mathbf{x}(t)$ provided an optimal schedule p .

Theorem 1. Let p be the optimal schedule for Problem 1. Thus, the optimal estimation for (1) at $t \in [\tau_k, \tau_{k+1})$ is given by the conditional mean $\mathbb{E}\{\mathbf{x}(t) \mid \mathbf{z}[0], \dots, \mathbf{z}[k-1]\}$ and can be computed as:

$$\begin{aligned} \hat{\mathbf{x}}(t) &= \mathbf{A}_d(t - \tau_k) \hat{\mathbf{x}}[k] \\ \hat{\mathbf{P}}(t) &= \text{cov}\{\mathbf{x}(t) - \hat{\mathbf{x}}(t)\} \\ &= \mathbf{A}_d(t - \tau_k) \hat{\mathbf{P}}[k] \mathbf{A}_d(t - \tau_k)^\top + \mathbf{W}_d(t - \tau_k) \end{aligned} \quad (5)$$

for $t \in [\tau_k, \tau_{k+1})$ with $\hat{\mathbf{x}}[k]$, $\hat{\mathbf{P}}[k]$ generated according to

$$\begin{aligned} \hat{\mathbf{x}}[0] &= \mathbf{x}_0, \quad \hat{\mathbf{P}}[0] = \mathbf{P}_0 \\ \mathbf{L}[k] &= \mathbf{A}_d(\Delta^{p_k}) \hat{\mathbf{P}}[k] \mathbf{C}^\top (\mathbf{C} \hat{\mathbf{P}}[k] \mathbf{C}^\top + \mathbf{R}^{p_k})^{-1} \\ \hat{\mathbf{x}}[k+1] &= \mathbf{A}_d(\Delta^{p_k}) \hat{\mathbf{x}}[k] + \mathbf{L}[k] (\mathbf{z}[k] - \mathbf{C} \hat{\mathbf{x}}[k]) \\ \hat{\mathbf{P}}[k+1] &= (\mathbf{A}_d(\Delta^{p_k}) - \mathbf{L}[k] \mathbf{C}) \hat{\mathbf{P}}[k] (\mathbf{A}_d(\Delta^{p_k}) - \mathbf{L}[k] \mathbf{C})^\top \\ &\quad + \mathbf{L}[k] \mathbf{R}^{p_k} \mathbf{L}[k]^\top + \mathbf{W}_d(\Delta^{p_k}) \end{aligned} \quad (6)$$

Proof. The proof can be found in Appendix C.1.

4. Perception latency scheduling

Now that the estimation stage of PLATE has been established, we turn our attention to obtaining the optimal scheduling policy, minimizing (2). In contrast to the estimations $\hat{\mathbf{x}}(t)$ which may take arbitrary values in \mathbb{R}^{n_x} , the perception schedule p is of discrete nature.

One possible way to obtain the optimal schedule that solves Problem 1 is to use Algorithm 2 by calling $\text{dynProg}(0, \hat{\mathbf{P}}[0], T_f)$. This algorithm is essentially an exhaustive search for the optimal schedule, organized as a dynamic programming recursive algorithm.

Fig. 1 illustrates how our dynamic programming algorithm works. The algorithm starts at time $\tau^+ = 0$ when calling $\text{dynProg}(0, \hat{\mathbf{P}}[0], T_f)$ and tries out each of the D possible perception methods. For each method, it computes a cost-to-arrive J (line 6 of Algorithm 2) and checks whether the current schedule length τ^+ is less than the desired length T_f . If $\tau^+ < T_f$, the algorithm must continue exploring other scheduling options that may extend the current schedule. This process is represented by the black nodes in Fig. 1. On the other hand, if $\tau^+ \geq T_f$, the algorithm has found a valid schedule that covers the entire time window $[0, T_f]$. This is represented by the white nodes in the figure.

At each step of the algorithm, a recursive call to the function is made using the current schedule length τ^+ and the updated covariance state $\hat{\mathbf{P}}[k+1]$. The recursion continues until the algorithm reaches the end of the time window T_f . At that point, the algorithm returns the optimal schedule decision p_k and the final cost for that schedule.

The correctness and complexity of the algorithm are established in the following result:

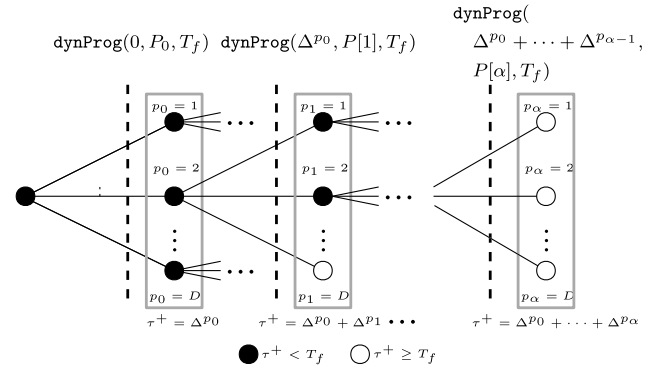


Fig. 1. Graphical depiction of Algorithm 2. At each step, all scheduling method options are explored in a recursive fashion with a maximum depth until the condition $\sum_{k=0}^{\alpha} \Delta p_k \geq T_f$ where $\alpha = \text{att}(p; [0, T_f])$ is reached.

Proposition 2. Calling $\{p, J\} \leftarrow \text{dynProg}(0, \hat{\mathbf{P}}[0], T_f)$ and computing $\hat{\mathbf{x}}(t)$ using the PLATE structure in (5) for such p , results in the optimal estimations, schedule and cost $\hat{\mathbf{x}}, p, \mathcal{J}(\hat{\mathbf{x}}, p) \equiv J$ for Problem 1 with worst case complexity given by $O(D^{\alpha_{\max}})$ where $\alpha_{\max} := \lfloor T_f / \min\{\Delta^1, \dots, \Delta^D\} \rfloor$.

Proof. The proof can be found in Appendix C.2.

Algorithm 2: dynProg

Data: $\tau, \hat{\mathbf{P}}[k], T_f$
Result: p, J

```

1  $J \leftarrow \infty$ 
2  $p \leftarrow \emptyset$ 
3 for  $\rho \in \{1, \dots, D\}$  do
4    $\tau^+ \leftarrow \tau + \Delta^\rho$ 
5    $p^+ \leftarrow \emptyset$ 
6    $J_\rho \leftarrow \frac{1}{T_f} \left( \lambda_\alpha r^\rho + \int_\tau^{\min(\tau^+, T_f)} \text{tr}(\mathbf{W}_d(t - \tau)) dt + \right.$ 
7      $\left. \int_\tau^{\min(\tau^+, T_f)} \text{tr}(\mathbf{A}_d(t - \tau) \hat{\mathbf{P}}[k] \mathbf{A}_d(t - \tau)^\top) dt \right)$ 
8   if  $\tau^+ < T_f$  then
9     Compute  $\hat{\mathbf{P}}[k+1]$  with (6) for  $\rho$  and  $\hat{\mathbf{P}}[k]$ 
10     $\{p^+, J^+\} \leftarrow \text{dynProg}(\tau^+, \hat{\mathbf{P}}[k+1], T_f)$ 
11     $J_\rho \leftarrow J_\rho + J^+$ 
12  end
13 if  $J_\rho < J$  then
14    $J \leftarrow J_\rho$ 
15   Append  $\rho$  to the start of  $p^+$  and assign the result to  $p$ 
16 end
```

The previous result evidences some important complications of Problem 1. First, as shown in 2, the complexity of the exact solution in Algorithm 2 increases exponentially as T_f increases. This is not surprising, since the discrete nature of the perception schedule p suggests that the problem is subject to a combinatorial explosion. Moreover, due to the transient behavior of $\hat{\mathbf{P}}(t)$, choosing p_k just to minimize a one step ahead of the cost in line 6 of dynProg may not be sufficient to find an exact solution for the problem without having to explore future scheduling method options over the window $[0, T_f]$. Some performance improvements can be made to Algorithm 2 such as applying a branch-and-bound technique [36, Chapter 2.3.3] or general approximations or heuristics [36, Chapter 6]. However, seeking for

tailor-made approximate solution and studying its performance gap with respect to the optimal is more appropriate for a practical implementation.

4.1. Quantized covariance approach

We provide an approximate alternative to Algorithm 2 in order to make it computationally feasible for use in PLATE. This approximation is sub-optimal, but can be made as close to optimal as desired, jeopardizing the computational complexity. The key idea is to observe that as the window $[0, T_f]$ is traversed during Algorithm 2, the set of possible states of the covariance $\hat{\mathbf{P}}[k]$ increase exponentially in size as well. This is evident from Fig. 1 which shows how the nodes at each level of the exploration tree increase exponentially as all possible combinations of perception methods are tested. However, if the amount of nodes at each level of the tree in Fig. 1 is managed to keep a bound regardless of T_f , the complexity of the algorithm can be reduced. This can be performed by grouping similar values of $\hat{\mathbf{P}}[k]$ into a single element of $\mathbb{R}^{n_x \times n_x}$ through quantization. The result is a compressed exploration graph with a non-growing number of nodes per level.

The quantization procedure for the covariance values is described as follows. Consider a compact set $\mathcal{B}_0 \subset \mathbb{R}^{n_x \times n_x}$ containing only positive semi-definite matrices, characterized by a bound $B_0 > 0$ such that $\|\hat{\mathbf{P}}[k]\|_F \leq B_0, \forall \hat{\mathbf{P}}[k] \in \mathcal{B}_0$. Now, given some $\delta > 0$, partition \mathcal{B}_0 into $Q_0(\delta) \in \mathbb{N}$ non-overlapping regions $\mathcal{B}_1, \dots, \mathcal{B}_{Q_0(\delta)}$ such that $\sup_{\hat{\mathbf{P}}[k], \hat{\mathbf{P}}'[k] \in \mathcal{B}_q} \|\hat{\mathbf{P}}[k] - \hat{\mathbf{P}}'[k]\|_F \leq \delta, \forall q \in \{1, \dots, Q_0(\delta)\}$ and $\bigcup_{q=1}^{Q_0(\delta)} \mathcal{B}_q = \mathcal{B}_0$. Moreover, pick a representative $\hat{\mathbf{P}}_q \in \mathcal{B}_q$ as an identifier for all other $\hat{\mathbf{P}}[k] \in \mathcal{B}_q$. Finally, let the quantization function $\mathcal{Q} : \mathcal{B}_0 \rightarrow \{\hat{\mathbf{P}}_1, \dots, \hat{\mathbf{P}}_{Q_0(\delta)}\}$ which takes any $\hat{\mathbf{P}}[k] \in \mathcal{B}_0$ and maps it to the $\hat{\mathbf{P}}_q$ such that $\hat{\mathbf{P}}[k] \in \mathcal{B}_q$. Note that we do not require to use a uniform quantization procedure. In fact, for practical purposes it is convenient to use a non-uniform quantization scheme in which $Q_0(\delta)$ points $\hat{\mathbf{P}}_1, \dots, \hat{\mathbf{P}}_{Q_0(\delta)} \in \mathcal{B}_0$ are provided instead, e.g. sampled from \mathcal{B}_0 , from which the maximum distance $\delta > 0$ is obtained.

The idea is to track how the identifiers $\hat{\mathbf{P}}_q$ with $q \in \{1, \dots, Q_0(\delta)\}$ are related between them using the PLATE estimation stage evolution in (6). To do so, a weighted directed graph $(\mathcal{V}, \mathcal{E}, \rho)$ is constructed with vertex set $\mathcal{V} = \{1, \dots, Q_0(\delta)\}$. Moreover, each edge $e = (i, j) \in \mathcal{E}$ corresponds to a connection between $i, j \in \{1, \dots, Q_0(\delta)\}$ when $\hat{\mathbf{P}}_j = \mathcal{Q}(\hat{\mathbf{P}}[k+1])$ with $\hat{\mathbf{P}}[k+1]$ computed using (6) for $\hat{\mathbf{P}}[k] = \hat{\mathbf{P}}_i$ and a weight function $\rho : \mathcal{E} \rightarrow \{1, \dots, D\}$ with $\rho(e) = p_k$.

However, it might be the case that $\hat{\mathbf{P}}[k+1]$ lies outside of \mathcal{B}_0 for some p_k and $\hat{\mathbf{P}}[k] = \hat{\mathbf{P}}_q, q \in \{1, \dots, Q_0(\delta)\}$. To account for these cases, we apply the steps described in Algorithm 3 adding new states outside \mathcal{B}_0 as required. The result is a graph \mathcal{G} with $Q(\delta) \geq Q_0(\delta)$ states where $\hat{\mathbf{P}}_q \in \mathcal{B}, \forall q \in \{1, \dots, Q(\delta)\}$ and \mathcal{B} is a region of $\mathbb{R}^{n_x \times n_x}$ with $\mathcal{B}_0 \subset \mathcal{B}$, and its corresponding bound $\|\hat{\mathbf{P}}[k]\|_F \leq B, \forall \hat{\mathbf{P}}[k] \in \mathcal{B}$. When Algorithm 3 finishes, it is ensured that if $\hat{\mathbf{P}}[0] \in \mathcal{B}_0$, then the quantized covariance trajectories will be contained in \mathcal{B} for any perception schedule.

It is beneficial for practical purposes, to estimate how $Q(\delta)$ grows with respect to original size of \mathcal{B}_0 as a result of Algorithm 3. The following result ensures that Algorithm 3 finishes with finite $Q(\delta)$. In addition, we provide an explicit upper bound for B , from which a worst case of $Q(\delta)$ can be computed depending on the actual quantization scheme.

Proposition 3. Algorithm 3 finishes for any compact set $\mathcal{B}_0 \subset \mathbb{R}^{n_x \times n_x}$ and $\delta > 0$. In addition, the resulting bound B for \mathcal{B} complies

$$B \leq B_s := \sqrt{n_x} \left(\frac{\lambda_{\max}(\Omega)}{\lambda_{\min}(\Omega)} \right) \left(B_0 + \frac{\bar{G}}{1 - \gamma} \right) \quad (7)$$

Algorithm 3: expandB

Data: δ, \mathcal{B}_0 with $\|\hat{\mathbf{P}}[0]\| \leq B_0$
Result: \mathcal{B}, \mathcal{G}

- 1 Quantize \mathcal{B}_0 into $Q_0(\delta)$ patches $\{\mathcal{B}_1, \dots, \mathcal{B}_{Q_0(\delta)}\}$ with identifiers $\{\hat{\mathbf{P}}_1, \dots, \hat{\mathbf{P}}_{Q_0(\delta)}\}$
- 2 $q \leftarrow 0$
- 3 $\mathcal{V} \leftarrow \{1, \dots, Q_0(\delta)\}$
- 4 $\mathcal{E} \leftarrow \emptyset$
- 5 $\mathcal{B} \leftarrow \mathcal{B}_0$
- 6 $Q(\delta) \leftarrow Q_0(\delta)$
- 7 **while** $q \leq Q(\delta)$ **do**
- 8 **for** $p \in \{1, \dots, D\}$ **do**
- 9 Compute $\hat{\mathbf{P}}[k+1]$ from (6) using $\hat{\mathbf{P}}[k] = \hat{\mathbf{P}}_q$ for $p_k = p$
- 10 **if** $\hat{\mathbf{P}}_{q'} \neq \mathcal{Q}(\hat{\mathbf{P}}[k+1])$ for any $q' \in \{1, \dots, Q(\delta)\}$ **then**
- 11 $Q(\delta) \leftarrow Q(\delta) + 1$
- 12 $\mathcal{V} \leftarrow \mathcal{V} \cup \{Q(\delta)\}$
- 13 Create a new patch $\mathcal{B}_{Q(\delta)}$ and identifier $\hat{\mathbf{P}}_{Q(\delta)}$ such that $\hat{\mathbf{P}}_{Q(\delta)}, \mathcal{Q}(\hat{\mathbf{P}}[k+1]) \in \mathcal{B}_{Q(\delta)}$ and $\sup_{\hat{\mathbf{P}}, \hat{\mathbf{P}}' \in \mathcal{B}_{Q(\delta)}} \|\hat{\mathbf{P}} - \hat{\mathbf{P}}'\| = \delta$
- 14 $\mathcal{B} \leftarrow \mathcal{B} \cup \mathcal{B}_{Q(\delta)}$
- 15 **end**
- 16 Add edge $e = (q, q')$ to \mathcal{E} with weight $\rho(e) = p$
- 17 **end**
- 18 $q \leftarrow q + 1$
- 19 **end**
- 20 Construct \mathcal{G} using \mathcal{V}, \mathcal{E} and its weights from line 16.

where $\Omega \in \mathbb{R}^{n_x \times n_x}$ and $\mathbf{Y}^i \in \mathbb{R}^{n_x \times n_z}, i \in \{1, \dots, D\}$ satisfy the following Linear Matrix Inequality (LMI):

$$\begin{bmatrix} \gamma \Omega & (\Omega \mathbf{A}_d(\Delta^i) - \mathbf{Y}^i \mathbf{C})^\top \\ (\Omega \mathbf{A}_d(\Delta^i) - \mathbf{Y}^i \mathbf{C}) & \Omega \end{bmatrix} \preceq 0, \quad (8)$$

$$i = 1, \dots, D, \quad \Omega \succ 0$$

for some $0 < \gamma < 1$. Moreover, $\lambda_{\min}(\Omega), \lambda_{\max}(\Omega)$ are the minimum and maximum eigenvalues of Ω respectively and $\bar{G} := \max_{i \in \{1, \dots, D\}} \|(\mathbf{L}^{p_k} \mathbf{R}^{p_k} (\mathbf{L}^{p_k})^\top + \mathbf{W}_d(\Delta^{p_k}))\|_F$ with $\mathbf{L}^{p_k} = \Omega^{-1} \mathbf{Y}^{p_k}$.

Proof. The proof can be found in Appendix C.3

Using the graph \mathcal{G} , the evolution of covariance states in \mathcal{V} can be tracked from an initial condition $q_0 \in \mathcal{V}$ as time advances in $[0, T_f]$. This can be visualized using a transition diagram, as shown in Fig. 2 which depicts an example with $D = 2$ and $\Delta^1 = \Delta_s, \Delta^2 = 2\Delta_s$. Here, at $t = 0$ the initial state is q_0 corresponding to $\hat{\mathbf{P}}_{q_0} = \mathcal{Q}(\mathbf{P}[0])$. We consider a discrete time $t = \ell \Delta_s$ as integer multiples of the minimum sampling interval. As ℓ increases, edges $e \in \mathcal{E}$ connect states between time steps, with different perception decisions $\rho(e)$. Note that, under this perspective, number of states at each time step is maintained constant. Hence, a dynamic programming algorithm for finite-state deterministic systems can be applied in order to obtain the optimal schedule p [36, Page 64]. In this context, the optimal p corresponds to the shortest route starting at q_0 , which traverses the transition diagram of \mathcal{G} until it reaches any state at stage $\ell = \lfloor T_f / \Delta_s \rfloor$ with distance measured by $\mathcal{J}(\hat{\mathbf{x}}, p)$.

The concrete steps required to run this dynamic programming solution, called qDP, are described in Algorithm 5 in Appendix A. Moreover, the sense in which we verify correctness of the dynamic programming as from Algorithm 5 and its complexity is established in the following:




Fig. 2. Transition graph example for covariance evolution using \mathcal{G} . In this example $D = 2$. Each node labeled in $1, \dots, Q(\delta)$ is connected to other nodes through one of the two perception decisions, either one or two steps ahead since $\Delta^1 = \Delta_s$, $\Delta^2 = 2\Delta_s$ for this example.

Proposition 4. Consider the qDP algorithm and assume that the evolution of $\hat{\mathbf{P}}[k]$ is constrained to evolve according to the structure in \mathcal{G} and that $\hat{\mathbf{P}}[0] \in \{\hat{\mathbf{P}}_1, \dots, \hat{\mathbf{P}}_{Q(\delta)}\}$. Thus, qDP obtains the optimal value of (2) for such constrained trajectories. Moreover, the worst-case of complexity of qDP is $O(\alpha_{\max} Q(\delta) D)$ where $\alpha_{\max} = \lfloor T_f / \Delta_s \rfloor$.

Proof. The proof can be found in Appendix C.4.

As evidenced by the previous result, the complexity of the quantized covariance approach is reduced to something linear in T_f . Nonetheless, since this is an approximate solution, there will be a trade-off between complexity and the performance gap of the quantized solution with respect to the true optimal. Note that by decreasing δ , the number of patches \mathcal{B}_q needed to cover \mathcal{B} will increase. However, it is expected that as $\delta \rightarrow 0$, the resulting sub-optimal solution performance improves. These ideas are formalized in the following result.

Theorem 2. Let \mathcal{J} be the optimal cost for Problem 1 and $\mathbf{P}[0] \in \mathcal{B}_0$. Then, for any $\varepsilon > 0$ there exists sufficiently small $\delta > 0$ such that

$$|\mathcal{J} - \mathcal{J}_Q(\delta)| \leq \varepsilon$$

where $\mathcal{J}_Q(\delta)$ is the cost obtained from qDP in Appendix A for such δ and initial condition $\mathcal{Q}(\hat{\mathbf{P}}[0]) \in \mathcal{B}_0$.

Proof. The proof can be found in Appendix C.5.

Henceforth, we have established that the PLATE strategy with estimation stage in (6) and perception-scheduling policy obtained using qDP is a sub-optimal for Problem 1, with mild computational complexity in terms of the window size T_f and can approximate the optimal solution with arbitrary precision if more computational power is available.

Remark 3. It is important to note that only the perception schedule computation is impacted by quantization, while the actual estimator in (5) and (6) is not quantized. Therefore, the quantization step δ only affects the resulting quality of the perception schedules in terms of the cost (2). We provide an analysis of the asymptotic effect of quantization on the cost in Theorem 2. While decreasing δ is expected to increase the quality of the perception schedules, it is not obvious whether the performance can be made arbitrarily close to the optimal one. This accuracy feature is ensured by Theorem 2.

Remark 4. Our approach can be extended to filter structures different than (6) such as a particle filter, as long as a similar

transition graph for the covariance matrix is obtained for its use in Algorithm 3. However, obtaining similar theoretical guarantees as in this work for other filter structures is not a trivial task and require more in-depth analysis.

5. Moving horizon PLATE

One property of the qDP method used in PLATE is that given an initial condition $\hat{\mathbf{P}}[0]$, a schedule p for the whole time window $[0, T_f]$ is obtained. Thus, it only suffices to compute the schedule at the beginning and traverse it element by element in line 5 of Algorithm 1 after each perception method is used. However, in a more practical TT scenario, there will be missing measurements as a result of occlusions, or distinguishability problems between targets in a multi-target setting. An appropriate re-detection of the targets may be performed by a maintenance mechanism as widely discussed in the literature. This may alter the quality of the pre-computed perception schedule.

In addition, the processed measurement quality, represented by the covariances $\{\mathbf{R}^1, \dots, \mathbf{R}^D\}$ may not remain fixed in a practical setting. Instead, it is usual to use the current estimation for the state of the target to improve the quality of subsequent measurements as in [37–39]. Moreover, the detection method may include some measure of uncertainty of their current output, e.g. through Bayesian techniques [40], or data-driven methods [41]. Thus, an online covariance estimate $\mathbf{R}[k]$ for the processed measurements may be available and may not be any of the nominal covariances. This discussion motivates to change the perception schedule according to the current state of the system by using qDP as a predictive policy.

To do so, we propose to use moving-horizon scheme for PLATE in the following way. First, we construct the transition graph \mathcal{G} as described in the previous section for nominal covariances $\{\mathbf{R}^1, \dots, \mathbf{R}^D\}$. Now, the dynamic programming algorithm qDP is used at τ_k with initial condition given by the current $q_0 = \mathcal{Q}(\hat{\mathbf{P}}[k])$ to obtain a perception schedule $p' = \text{qDP}(q_0, T_f, \mathcal{G})$ for the next window $[\tau_k, \tau_k + T_f]$. As a result, we apply the first perception method of p' during $[\tau_k, \tau_k + \Delta^{p_k})$ and repeat the procedure for the next interval.

Furthermore, every time a new perception output is obtained, the PLATE correction in (6) is computed with the actual processed measurement uncertainty $\mathbf{R}[k]$ if available, otherwise with its nominal covariance. When there are no measurements, we simply predict the state of the target and its covariance using (5). This procedure is summarized in Algorithm 4 which describes the steps that must be performed in line 5 of Algorithm 1 in order to implement this strategy.

Moreover, the latency of executing qDP at each τ_k can be considered negligible if two of the following approaches is used. First, since $\hat{\mathbf{P}}[k]$ is quantized through $\mathcal{Q}(\bullet)$ there are only $Q(\delta)$ possible outcomes for which their resulting perception method obtained through qDP can be pre-computed offline. On the other hand, similarly to what is suggested in [22], taking $\hat{\mathbf{P}}[k]$ and converting it to a perception method in $\{1, \dots, D\}$ is a classification problem with a low dimensional input. Thus, the perception decision can be learned by a classifier.

Remark 5. Comparing our approach with previous work, note that the perception scheduling policy as in Algorithm 4 acts as a frame-skipping technique similar to [1,18,19]. However, unlike prior frame-skipping approaches which are based on heuristic rules, we provide an actual performance guarantees for our sub-optimal perception schedules as given in Theorem 2. On the other hand, note that unlike related multi-sensor scheduling approaches in the literature as in [26–32], Algorithm 5 can be extended for different cost function expressions by modifying line 12 of Algorithm 6, which makes our proposal more versatile.

Algorithm 4: Moving-horizon PLATE

```

// Offline:
1 Pre-compute  $\mathcal{G}$  using Algorithm 3, alternatively
  pre-compute  $\text{qDP}(\mathcal{Q}(\hat{\mathbf{P}}_q, \mathcal{G}, T_f)$  for all
   $\hat{\mathbf{P}}_q, q \in \{1, \dots, Q(\delta)\}$ .
// Online:
2 if There is new processed measurement then
3   Update  $\hat{\mathbf{P}}[k]$  using (6) for the previous scheduling
  decision  $p_{k-1}$  and its resulting perception uncertainty
   $\mathbf{R}[k-1]$ .
4 else
5   Predict  $\hat{\mathbf{P}}[k]$  using (5) for the previous scheduling
  decision  $p_{k-1}$ .
6 end
7  $p', \_ \leftarrow \text{qDP}(\mathcal{Q}(\hat{\mathbf{P}}[k]), \mathcal{G}, T_f)$ 
8  $p_k \leftarrow$  first element of  $p'$ 

```

Remark 6. The moving horizon PLATE proposal allows us to change the system parameters online as in line 3 of Algorithm 4, where we use an uncertainty $\mathbf{R}[k-1]$ which need not to be any of the nominal ones for the perception methods. This poses an important advantage with respect to the philosophy of some multi-sensor scheduling works such as [26,29] which rely on periodic schedules or other strategies that are highly dependant on the problem structure.

6. Numerical experiments

In order to evaluate PLATE we consider the following scenario. First, let a target with state $\mathbf{x} = [x, v_x, y, v_y]^T$ describing its position and velocity on the plane. Let camera images available every $\Delta_s = 1/30$ seconds, coinciding with the usual frame rate of a camera. In addition, the position $[x, y]^T$ is measured through a detection process from which there are $D = 2$ available perception configuration. First, a lightweight detector with latency $\Delta^1 = 3\Delta_s$ and nominal covariance $\mathbf{R}^1 = \text{diag}(0.5, 0.5)$. Second, a more accurate detector with latency $\Delta^2 = 9\Delta_s$ and nominal covariance $\mathbf{R}^2 = \text{diag}(0.05, 0.05)$. Motivated by Remark 2 we consider CPU loads of $f^1 = 0.5$ and $f^2 = 0.8$ resulting in penalties $r^1 = 0.5\Delta^1$ and $r^2 = 0.8\Delta^2$. The system matrices in this scenario are:

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

In addition, let the process covariance $\mathbf{W} = \text{diag}(0.5, 0.5)$. This double integrator target model is often used for generic targets in the literature. In the following, we evaluate the different aspects of our proposals for this setting.

6.1. Numerical covariance bound estimation

We start by building a transition graph \mathcal{G} . To do so, we set $\mathcal{B}_0 \in \mathbb{R}^{n_x \times n_x}$ as the set of all positive definite matrices \mathbf{P} with $\|\mathbf{P}\|_F \leq \mathbf{B}_0 := 1$ and use the mechanism in Algorithm 3. In order to depict the result in Proposition 3, we explicitly obtain a bound for \mathcal{B} . First, note that the LMI (8) has a solution for $\gamma = 0.98$ as

$$\mathbf{Y}^1 = \begin{bmatrix} 0.031 & 0 \\ 0.037 & 0 \\ 0 & 0.031 \\ 0 & 0.037 \end{bmatrix}, \mathbf{Y}^2 = \begin{bmatrix} 0.122 & 0 \\ 0.137 & 0 \\ 0 & 0.122 \\ 0 & 0.137 \end{bmatrix}$$




Fig. 3. Covariance norm evolution, showing the bound B_s as well as $\|\hat{\mathbf{P}}(\tau_k)\|_F$ for PLATE with 100 random initial conditions $\hat{\mathbf{P}}[0]$ and schedules p with $\|\hat{\mathbf{P}}[0]\|_F \leq B_0 = 1$.

with a resulting bound (7) of $B_s = 4.922$. Fig. 3 depicts how given initial conditions $\mathbf{P}[0] \in \mathcal{B}_0$ and a randomly generated perception schedule, the magnitude of the PLATE covariance complies $\|\hat{\mathbf{P}}[k]\|_F \leq B_s, \forall k \geq 0$.

6.2. Cost comparison using qDP

Now, in order to depict Theorem 2 through this example, we use different levels of quantization for evaluation. Note that the results presented until now do not require uniform quantization. In fact, in the following examples for simplicity we use non-uniform quantization by sampling a fixed number of points $\hat{\mathbf{P}}_1, \dots, \hat{\mathbf{P}}_{Q(\delta)} \in \mathcal{B}_0$ and quantizing any other $\hat{\mathbf{P}} \in \mathcal{B}_0$ through a nearest neighbor rule under the $\|\bullet\|_F$ norm. Next, we expand \mathcal{G} through Algorithm 3 to obtain \mathcal{B} . We use $Q(\delta) = 50, 500, 5000$ which after this process results in $\delta = 10.45, 4.23, 2.14$ respectively for these arbitrarily selected points, drawn randomly over \mathcal{B}_0 with uniform distribution, in order to analyze the impact of the quantization step in the tracking quality.

We evaluate qDP for 100 random samples for $\mathbf{P}[0]$ over \mathcal{B} and compute their resulting cost (2). First, consider $T_f = 1, \lambda_\alpha = 5$. In this case, it is possible to obtain the minimum cost J_{\min} by evaluating all schedules covering the window $[0, T_f]$. We show the results in the first row of Fig. 4 where the histograms depict the frequency of the distance $|J_{\min} - J_p|$ where J_p is the cost obtained for a schedule p . The schedules tested were obtained through qDP for graphs \mathcal{G} using the previously described values of $Q(\delta)$, as well as the static schedules $p = \{1, 1, \dots\}$ and $p = \{2, 2, \dots\}$. The results show that as the quantization gets finer, the cost concentrate more and more towards the minimum, i.e. $|J_{\min} - J_p|$ becomes smaller. In addition, it is observed that even with a coarse discretization $Q(\delta) = 50$, the cost of qDP is almost always less than the cost for the static schedules. Moreover, the average CPU load seems to decrease as well as the discretization gets finer. A similar behavior is obtained when using $T_f = 10, \lambda_\alpha = 100$, as depicted in the second row of Fig. 4. In this case, J_{\min} is computed for 10000 different randomly selected perception schedules aiming to approximate the otherwise intractable exhaustive search for the true optimal cost.

Another interesting experiment is obtained by increasing $\lambda_\alpha = 15$ such that the cost of the covariance is negligible to the one of the penalties r^{p_k} in (2). In this case, J_{\min} is directly the static schedule $p = \{1, 1, \dots\}$ which is the schedule with less CPU load. In addition, since the particular values of covariance $\hat{\mathbf{P}}[k]$ have little impact on the final cost, it only suffices to use a single point $Q(\delta) = 1$ in the discretization of the covariance space. Hence, the result of qDP is the static schedule $p = \{1, 1, \dots\}$ regardless of the initial condition $\hat{\mathbf{P}}[0]$ as shown in the third row of Fig. 4.

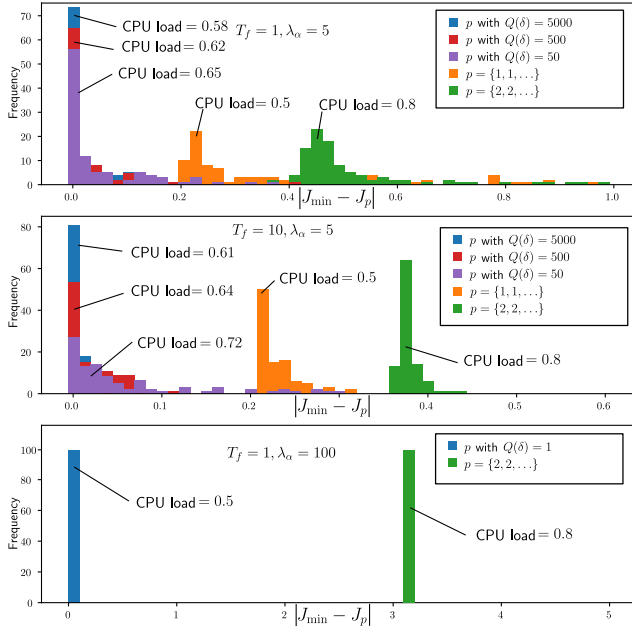


Fig. 4. Resulting histograms for the cost difference $|J_{\min} - J_p|$ with the schedule obtained from the approximate dynamic programming approach for 100 random initial conditions for $P[0]$ as described in Section 6.2. The parameters T, λ_α of the cost function in (2) were changed for three different scenarios as well as the number of quantization points $Q(\delta)$. In addition, only the average CPU load for all experiments is shown for convenience.

6.3. Moving-horizon PLATE scheduling

We simulate (1) for $\mathbf{x}_0 = 0, P[0] = 4I$ using the Euler-Maruyama discretization with time step $\Delta t = 10^{-3}$ and run the loop in Algorithm 1 under the moving-horizon scheduling of Algorithm 4 where qDP is configured similarly as in the previous examples, with $T_f = 10, \lambda_\alpha = 5, Q(\delta) = 5000, B_0 = 5$ (Algorithm 3). Fig. 5 depicts the time evolution of $\mathbf{x}(t)$ as well as the PLATE estimate $\hat{\mathbf{x}}(t)$, where we show only x and v_x for convenience, since y and v_y behave similarly. In addition, we show confidence intervals of 3 times the standard deviation for each coordinate obtained from $\hat{P}(t)$. In addition, we show the scheduling decision p_k at each time. In order to compare the cost in this experiment we show $\text{tr}(\hat{P}(t))$ for the moving-horizon scheduling, as well as for the static schedules. It is worth noting that the best quality measured with $\text{tr}(\hat{P}(t))$ is obtained by the static schedule with $p_k = 2$. However, it has the highest CPU load of 0.8. The moving-horizon PLATE strategy manages to have an intermediate quality between the two static schedules, with a CPU load of 0.65 for this experiment. Thus, PLATE manages to obtain a better trade-off between quality and resource usage. In addition, an occlusion is simulated in the interval $t \in [4, 6]$ which shows how uncertainty increases during this period, but recovers once new measurements arrive.

7. Evaluation on real data

In this section, we evaluate PLATE on a real target tracking task. We use the MOT16 benchmark from [42] since it is part of a public standard data-set, widely used to evaluate target tracking algorithms. In the simplest setting, the task is to take images from a monocular camera and track a target of interest across multiple frames. As usual, and for the sake of simplicity, we are interested in tracking the pixel position of the geometric center of a bounding box surrounding the target. For this example, we

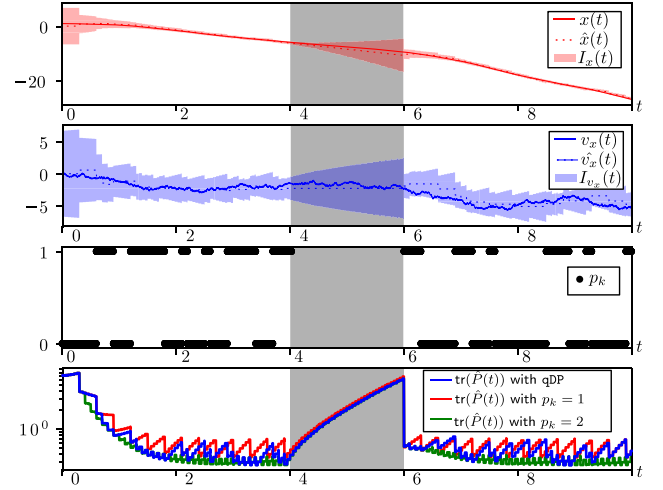


Fig. 5. Behavior of the target model (1) and the output of PLATE using the moving-horizon scheduling with $T_f = 10$ and $Q(\delta) = 5000$ when an occlusion occurs for $t \in [4, 6]$. Moreover, $I_x(t)$ and $I_{v_x}(t)$ represent confidence intervals around $\hat{x}(t)$ and $\hat{v}_x(t)$ of 3 times their standard deviation.

model the pixel position with a two-dimensional single integrator with fixed $n_x = n_z = n_w = 2, \mathbf{A} = \mathbf{0} \in \mathbb{R}^{2 \times 2}, \mathbf{B} = \mathbf{C} = \mathbf{I} \in \mathbb{R}^{2 \times 2}$.

In the following, we describe a target tracking pipeline that consists of a perception stage composed of three different target detection methods based on neural networks, each exhibiting varying latency and accuracy performance, and a scheduling stage responsible for deciding which of these methods to use at each time. The objective of this section is to demonstrate that PLATE is the best choice for the scheduling stage. For this purpose, we compare the tracking performance of PLATE with two other approaches: (1) a standard neural network-based method, which involves using the perception methods separately with no scheduling, (2) event-triggered approaches for frame-skipping.

7.1. Evaluation framework

The MOT16 data-set includes 7 video sequences, each with labeled bounding-boxes for different moving targets. We manually selected several targets for each sequence. We recorded their ground-truth positions and the corresponding portion of the sequence, which we refer to as target tracks from now on. We chose targets that provided longer experiments, in order to obtain more comprehensive results. Our evaluation data consists of 50 target tracks across all 7 video sequences, whose ground-truth IDs are shown in Table 1 as a reference. Note that for all video sequences the frame rate is set to 30 frames per second except for MOT16-05 and MOT16-13 with frame rates of 14 and 25 frames per second respectively. As a result, target tracks for MOT16-13 and MOT16-05 (15 tracks in total) were used exclusively to estimate the nominal detection covariance for each perception method as described later, and the covariance \mathbf{W} in (1). This process is called training in Table 1. The rest of the 35 target tracks contained in the remaining 5 video sequences were used to evaluate the online performance of PLATE. From this point onward, all experiments (both for our proposal and other proposals used for comparison) were conducted on a PC equipped with an Intel Core i7-8700, along with the aid of an NVIDIA GeForce GTX 1080 Ti graphics card.

To evaluate the performance of PLATE, it is necessary to use a setup that includes a bank of perception methods. However, the purpose of our work is not to evaluate the individual

Table 1

Ground-truth target IDs, as provided by the MOT16 data-set, picked for each video sequence in our evaluation. Frames per second is abbreviated as fps. The last column indicates T or E if the target tracks in the corresponding video sequence are used for training or for evaluation respectively.

Name	Target IDs	fps	Usage
MOT16-05	2,37,41,124,128	14	T
MOT16-13	13,14,29,30,76,77,79,81,95,124	25	T
MOT16-02	3,19,20,32,33,34,35	30	E
MOT16-04	2,67,70,80,83,98,99,100	30	E
MOT16-09	12,13,19,20,23	30	E
MOT16-10	1,2,4,7,18,24	30	E
MOT16-11	1,3,4,5,6,11,25,27,29	30	E

performance of these methods, but rather to measure the advantages of PLATE in terms of target state estimation quality and resource usage, such as CPU load. Therefore, we selected well-established perception methods from the public model zoo of Detectron2 [43]. This repository provides many neural network-based target detection algorithms with varying levels of latency and precision, making it suitable for our work. Specifically, we chose the `faster_rcnn_R_50_FPN_3x` and `faster_rcnn_R_101_C4_3x` models (hereinafter referred as the fast and slow networks respectively), which represent the clearest examples in the zoo of low quality detection with short latency and good quality detection with long latency, respectively. These models have a reported latency of 0.038s and 0.104s on a standard computing platform, which were verified in our own setup. To further simplify the experiments, for the fast network we sub-sample the input image in a factor of two, reducing the latency to roughly $1/(30\text{fps}) \approx 0.033$ s which is the inverse of the frame rate. With these two networks we propose 3 perception methods as follows:

- $p_k = 1$: This method uses the fast network with a latency of $1/(30\text{fps})$. This means that this method does not skip any frame and has a CPU load of 100%.
- $p_k = 2$: It uses the slow network with a processing latency of 0.104s. This means that the network is computed over four frames, causing it to skip three frames, processing only the first one. Consequently, the CPU load over these four frames is $\frac{0.104\text{s}}{4\text{frames}/(30\text{fps})} = 78\%$.
- $p_k = 3$: In order to have an option which favors a low CPU load, we use the fast network, but deliberately skip the next 4 frames. This means that the perception quality of this method is the same as $p_k = 1$ but with a CPU load of $\frac{1/30\text{ s}}{5\text{frames}/(30\text{fps})} = 20\%$ and longer overall latency.

These perception method options are depicted graphically in Fig. 6. The previous framework has a disadvantage in terms of memory consumption, as the adaptive strategy requires storing both neural networks, leading to higher memory usage compared to using the perception methods individually. However, it is worth noting that depending on the application, it may be possible to use a single ANN, as demonstrated in [20], which can offer different latency and quality levels with a fixed amount of memory space.

As depicted in Fig. 7, the output of each of the neural networks is a list of bounding boxes for all the detected objects in the image. Therefore, a data association stage is necessary to select the appropriate bounding box for the target object in each experiment. There exist many data association techniques in the literature, such as those based on similarity measures using neural networks [44] or end-to-end detection and association siamese networks [45]. These methods are highly dependent on the training procedure, which might have an undesired effect

when evaluating the scheduling feature of PLATE by itself. Hence, since the goal of this work is to evaluate PLATE for a fixed perception framework, we use a more standard approach. We compare the predicted position of the target, $\lim_{t \rightarrow \tau_k} \hat{\mathbf{x}}(t)$, from (5) at the time τ_k when the image was captured, and use the Hungarian algorithm to select the most suitable bounding box candidate [1]. The computational latency of this procedure is negligible and can be ignored for simplicity in subsequent discussions. Fig. 7 shows that the output of the data association stage is the best bounding box candidate for the target of interest from which a processed position measurement $\mathbf{z}[k]$ is obtained as its geometric center.

In this setting, $\Delta_s = 1/30$, $\Delta^1 = \Delta_s$, $\Delta^2 = 3\Delta_s$, $\Delta^3 = 5\Delta_s$. In addition, the CPU loads are $f^1 = 1$, $f^2 = 0.78$, $f^3 = 0.2$. We estimate the nominal error covariance for each perception method $\mathbf{R}^1, \mathbf{R}^2, \mathbf{R}^3$ by comparing the outputs of each perception method with the ground-truth data across all images in the 15 training target tracks. The resulting covariance matrices are

$$\mathbf{R}^1 = \text{diag}(13.12^2, 25.87^2), \mathbf{R}^2 = \text{diag}(9.94^2, 17.06^2), \mathbf{R}^3 = \mathbf{R}^1$$

The matrix \mathbf{W} was estimated using the training target tracks as well, following a standard parameter estimation procedure described in [41]. As described in Section 5, we also aim to evaluate PLATE when an online estimation of the current covariance $\mathbf{R}[k]$ is available besides the nominal covariances. As a result, we adapt the standard parameter estimation in [41] for this setting, using the resulting performance of previous detections compared to the estimation $\hat{\mathbf{x}}(t)$. Let $\mathcal{I}^{p_k}[k] = \{\ell \leq k : p_\ell = p_k, k - \ell \leq N_w\}$ be the set of all discrete instants ℓ corresponding to the moments τ_ℓ in which the perception method p_k was used prior to the current τ_k , in a moving window picked here of $N_w = 10$ samples. Thus, if p_k is to be picked at $t = \tau_k$ one can estimate the current covariance for such method as:

$$\mathbf{R}[k] = \frac{1}{N_w} \sum_{\ell \in \mathcal{I}^{p_k}[k]} \mathbf{e}[\ell] \mathbf{e}[\ell]^T - \hat{\mathbf{C}}\mathbf{P}^-(\tau_k) \mathbf{C}^T \quad (9)$$

where $\mathbf{e}[\ell] = \mathbf{C}\hat{\mathbf{x}}[\ell] - \mathbf{z}[\ell]$ and $\hat{\mathbf{P}}^-(\tau_k) = \lim_{t \rightarrow \tau^-} \hat{\mathbf{P}}(t)$ from (5).

The PLATE module, depicted in Fig. 7, executes the estimator-predictor equations specified in (5) and (6), and implements the scheduler algorithm presented in Algorithm 4. This module employs a moving-horizon PLATE with $T_f = 10$ s and utilizes pre-computed schedules for each quantized covariance value $\hat{\mathbf{P}}_q$, which are based on nominal $\mathbf{R}^1, \mathbf{R}^2$, and \mathbf{R}^3 . When a new processed measurement $\mathbf{z}[k]$ is obtained, the pre-computed values for line 7 in Algorithm 4 for all possible $\hat{\mathbf{P}}[k]$ in a compact set enable the computation of the new scheduling decision p_k with negligible computing latency. Additionally, the computation time required for (5) and (6) in this example's state space dimensions is negligible compared to the latencies Δ^1, Δ^2 , and Δ^3 . The values of $Q(\delta)$ will be varied between 50, 500, 5000 as in previous examples with \mathcal{B}_0 as the set of all positive definite matrices \mathbf{P} with $\|\mathbf{P}\|_F \leq \mathbf{B}_0 = 20$.

Remark 7. The cost function (2) was set using $\lambda_\alpha = 1$ and

$$r^{p_k} = \lambda_{\text{load}} f^{p_k} \Delta^{p_k} + \lambda_{\text{att}}. \quad (10)$$

The term $\lambda_{\text{load}} f^{p_k} \Delta^{p_k}$ penalizes the CPU load, while λ_{att} penalizes the attention, weighted by $\lambda_{\text{load}}, \lambda_{\text{att}}$. Minimizing these objectives while providing the best possible accuracy are conflicting goals, but can be balanced by adjusting $\lambda_{\text{load}}, \lambda_{\text{att}}$. If $\lambda_{\text{load}} = \lambda_{\text{att}} = 0$, the offline construction of p_k results in $p_k = 1$ for every $\hat{\mathbf{P}}_q$, maximizing the attention regardless of the quantization step. Conversely, we verified numerically that for $\lambda_{\text{load}} = 1$ and $\lambda_{\text{att}} > 15$, $p_k = 3$ for all $\hat{\mathbf{P}}_q$, minimizing the attention, regardless of the quantization step as well. For other values of $\lambda_{\text{load}}, \lambda_{\text{att}}$, the schedule p_k changes according to the current $\hat{\mathbf{P}}_q$, achieving a balance between these objectives.

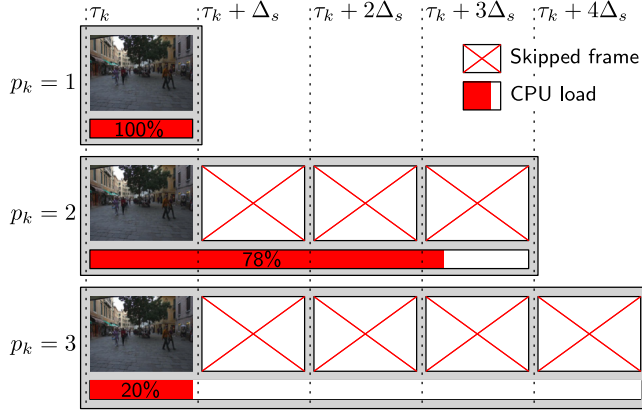


Fig. 6. Different features of the three perception methods described in Section 7.1. The number of skipped frames as well as the CPU load in each method is depicted.

7.2. Performance of the implemented pipeline

In this section, we evaluate the proposed pipeline on the 35 evaluation target tracks. The performance evaluation metric used is the Mean Squared Error (MSE) of $\hat{\mathbf{x}}(k\Delta_s)$, $k = 0, 1, \dots$ compared to the ground-truth data available in the MOT16 dataset. Additionally, the CPU load for all experiments is recorded, as well as the attention measured as the percentage of non-skipped frames compared to the total number of frames. To illustrate the trade-off between the previously described metrics, a combined cost is computed as $(\text{MSE}) + \lambda_{\text{load}}(\text{CPU load}) + \lambda_{\text{att}}(\text{Attention})$ as an sampled version of the cost (2) under penalties (10).

In the following, we used $\lambda_{\text{load}} = \lambda_{\text{att}} = 1/2$ in (10) which results in a trade-off between accuracy, CPU load and attention, producing different p_k through time for PLATE. The results are shown in Table 2 where we compare the performance of the proposal under different values of $Q(\delta)$ as well as using nominal $\mathbf{R}[k] \in \{\mathbf{R}^1, \mathbf{R}^2, \mathbf{R}^3\}$ or adaptive $\mathbf{R}[k]$ in (9).

There are three main conclusions that can be obtained from these results. First, note that using adaptive $\mathbf{R}[k]$ produces an improvement with respect to the nominal case, particularly with low $Q(\delta)$ where an improvement of $(35.76 - 32.78)/35.76 \approx 8\%$ in MSE is obtained when $Q(\delta) = 50$. Second, consistent with Theorem 2, increasing $Q(\delta)$ improves the MSE as well, with diminishing returns as $Q(\delta)$ increases. Third, the CPU load and attention remain roughly the same meaning that high values of $Q(\delta)$ might not be needed in practice to maintain reasonable values for these objectives.

Now, we compare PLATE with other ideas in the literature. As a baseline for the comparison we used a framework which does not use scheduling. This is, the neural network method to be used for perception is chosen at the beginning, and remains fixed at all times. The results are shown in the first three rows of Table 3. Note that the configuration in row (1) of Table 3 (fixed $p_k = 1$) obtains the best MSE among all experiments at the expense of high CPU load and attention. In addition, row (3) of Table 3 (fixed $p_k = 3$) produces the best results in terms of CPU load and attention at the expense of high MSE. However, recall from Remark 7 that with appropriate configuration of the penalties (10), PLATE can produce either fixed $p_k = 1$ or $p_k = 3$ and obtain the best performance for MSE, CPU load or attention by separate.

Moreover, we evaluated frame-skipping techniques [1,18,19] which rely on an event-triggered condition to determine whether a new frame needs processing. To evaluate a similar approach, we decided at each new frame whether to use $p_k = 1$ or skip the

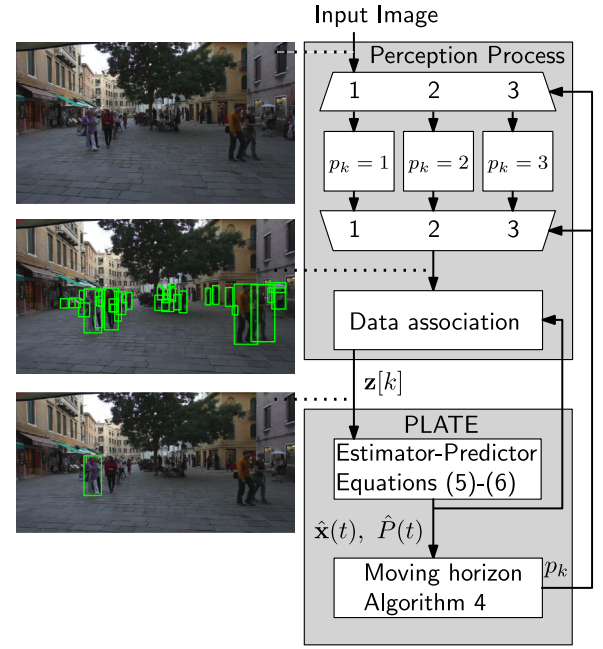


Fig. 7. Implemented Pipeline for the PLATE evaluation framework described in Section 7.1.

frame entirely. Send-on-delta strategies are commonly employed, where a frame is processed if $\text{tr}(\hat{\mathbf{P}}(k\Delta_s)) \geq \delta_{\text{ET}}$ with a threshold $\delta_{\text{ET}} > 0$. The results for different δ_{ET} are shown in rows (4) and (5) of Table 3.

Finally, row 6 of Table 3 shows the best configuration for PLATE as obtained in Table 2. For comparison, note that while fixed $p_k = 2$ yields a small MSE, CPU load and attention values remain high. In contrast, PLATE in row 6 of Table 3 also has a small MSE with the advantage of reducing the CPU load from 78% to 43.5% and attention from 75% to 43.4%, when compared to fixing $p_k = 2$.

The event-triggered approach in row 4 of Table 3 shows a small improvement in the MSE can be obtained when compared to PLATE in row 6 of Table 3. Despite this, the CPU load and attention values are considerably bigger for the event triggered alternative. More concretely, PLATE in row 6 of Table 3 improves the CPU load from 71.2% to 43.5% (a relative improvement of 38.9%) when compared to the event triggered approach in row 4 of Table 3, even with a similar MSE performance. This trade-off is clear since PLATE obtains the best performance among all options as illustrated with the combined cost in the last column of Table 3.

Hence, using PLATE with appropriate configuration in (10) can obtain a good trade-off between accuracy, CPU load and attention, with improved performance with respect to static perception configurations. In addition, PLATE provides a clear connection via the cost function (2) and the penalty (10) to achieve in each situation different behaviors, prioritizing CPU load and attention or MSE.

Remark 8. The performance of the proposal is highly dependent on the parameters $Q(\delta)$, λ_{load} , and λ_{att} , as evident from the results presented in Tables 2 and 3. It is important to note that the optimal parameter selection generally varies based on the specific application and target dynamics under consideration. However, as discussed in this section, conducting several experimental tests can help determine the most suitable parameters for the user.

Table 2

Evaluation of the implemented pipeline for PLATE on MOT16 data as described in Section 7. MSE stands for Mean-Squared-Error with respect to ground-truth data. The attention column corresponds to the percentage of non-skipped frames with respect to the total number of frames. The penalty in (10) is configured with $\lambda_{\text{load}} = \lambda_{\text{att}} = 1/2$. Moreover, to illustrate the trade-off between the previously described metrics, a combined cost is computed as $(\text{MSE}) + \lambda_{\text{load}}(\text{CPU load}) + \lambda_{\text{att}}(\text{Attention})$ as an sampled version of the cost (2) under penalties (10).

	$\mathbf{R}[k]$	$Q(\delta)$	MSE [px]	CPU load [%]	Attention [%]	Combined cost
(1)	Nominal	50	35.76	44.5	44.1	80.06
(2)	Adaptive	50	32.78	44.0	43.9	76.73
(3)	Nominal	500	32.14	44.1	43.9	76.14
(4)	Adaptive	500	32.02	43.8	43.6	75.72
(5)	Nominal	5000	31.54	43.9	43.6	75.29
(6)	Adaptive	5000	30.04	43.5	43.4	73.49

Table 3

Evaluation of the implemented pipeline. The evaluation metrics are the MSE, CPU load and attention as well as the combined cost taking into account all previous metrics weighted by $\lambda_{\text{load}} = \lambda_{\text{att}} = 1/2$. The methods under evaluation are: 1) No scheduling with fixed $p_k = 1$, corresponding to using only the fast neural network `faster_rcnn_R_50_FPN_3` at all times without additional frame skipping. 2) No scheduling with fixed $p_k = 2$ corresponding to using only the slow neural network `faster_rcnn_R_101_C4_3x` at all times without additional frame skipping. 3) No scheduling with fixed $p_k = 3$ corresponding to using only the fast neural network `faster_rcnn_R_50_FPN_3` followed by 4 skipped frames repeatedly. 4) and 5) Event triggered approach with different δ_{ET} . 6) The best configuration for PLATE as obtained from Table 2.

	Method	Configuration	MSE [px]	CPU load [%]	Attention [%]	Combined cost
(1)	No scheduling	fixed $p_k = 1$	19.61	100	100	119.61
(2)	No scheduling	fixed $p_k = 2$	29.05	78	75	105.55
(3)	No scheduling	fixed $p_k = 3$	73.73	20	20	93.73
(4)	Event-triggered	$\delta_{\text{ET}} = 5$	25.04	71.2	69.2	95.24
(5)	Event-triggered	$\delta_{\text{ET}} = 50$	64.74	35.9	35.2	100.29
(6)	PLATE	$Q(\delta) = 5000$, Adaptive $\mathbf{R}[k]$	30.04	43.5	43.4	73.49

Our experiments demonstrate a range of potential performance values obtained by varying the parameter $Q(\delta)$, allowing for selection based on specific application requirements and resource availability. In addition, regarding the penalties λ_{load} , and λ_{att} , these values should reflect the application specifications. Still, it is beneficial to identify parameter sets for these penalties that exhibit similar behavior to the ones described in Remark 7. This knowledge becomes particularly valuable when the user needs to make a trade-off between accuracy and resource usage. More principled design rules are not trivial to obtain and will be considered in our future work.

8. Conclusions

In this work, we introduced PLATE as a perception-latency aware estimator for target tracking applications. PLATE leverages a bank of perception methods with different latency-precision trade-offs to adaptively select the best method for the current estimation task. The proposed algorithm allows for the skipping of input frames, which reduces CPU load and resource usage while still maintaining high tracking accuracy. Unlike other frame skipping techniques, PLATE's algorithm is based on a formal dynamic programming argument rather than heuristics. We found that while the exact solution of the problem is subject to a combinatorial explosion, an approximate solution can be obtained with efficient computational complexity. We evaluated PLATE using both simulations and real-world data sets, and found that it outperforms other state-of-the-art approaches in terms of both tracking accuracy and computational efficiency. Our current approach is limited to linear target motion models, and extending the estimator to nonlinear models using particle filtering or data-driven techniques is challenging and left for future work.

Declaration of competing interest

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

Appendix A. Quantized covariance method

The concrete steps to perform the dynamic programming algorithm described in Section 4.1 are shown in Algorithm 5. First, auxiliary matrices \mathbf{M}_Q , \mathbf{M}_P , \mathbf{M}_J are computed from Algorithm 6 given an initial state q_0 . In line 6 of Algorithm 6 time steps up to α_{max} are traversed. Moreover, in line 7 all possible states q at time ℓ are evaluated as well. Furthermore, for each of these states q , in line 8 all perception options ρ are checked. In this way, Algorithm 6 tracks optimal routes where $[\mathbf{M}_J]_{q,\ell}$ stores the best cost-to-arrive from q_0 to q after ℓ time steps. Moreover, $[\mathbf{M}_Q]_{q',\ell}$ stores the best state q connected to q' at time step ℓ , where the perception decision is stored in $[\mathbf{M}_P]_{q',\ell}$ defining how many stages separate both states as well. These matrices are used in Algorithm 5 to trace-back the optimal scheduling.

Appendix B. Auxiliary results

Proposition 5. Let $\lambda = [\lambda_1, \dots, \lambda_n]^\top \in \mathbb{R}^n$ and define $\|\lambda\|_p = (\sum_{i=1}^n |\lambda_i|^p)^{1/p}$. Then, with $0 < r < s$, the following inequalities are satisfied:

- (a) [46, Theorem 16, Page 26] $\|\lambda\|_r \leq n^{\frac{1}{r} - \frac{1}{s}} \|\lambda\|_s$.
- (b) [46, Theorem 19, Page 28] $\|\lambda\|_s \leq \|\lambda\|_r$.

Algorithm 5: qDP

Data: q_0, T_f, \mathcal{G} computed from Algorithm 3
Result: p, J

- 1 $\{\mathbf{M}_Q, \mathbf{M}_p, \mathbf{M}_j\} \leftarrow \text{qDPMatrices}(q_0, T_f, \mathcal{G})$
- 2 Set J and q to the best cost and state at the last column of \mathbf{M}_j
- 3 $p \leftarrow \emptyset$
 // Use \mathbf{M}_p to trace back the optimal schedule ending at state q
- 4 $\ell \leftarrow \lfloor T_f / \Delta_s \rfloor$
- 5 **while** $\ell > 0$ **do**
- 6 $\rho \leftarrow \lfloor \mathbf{M}_p \rfloor_{q, \ell}$
- 7 Append ρ at the start of p
- 8 $q \leftarrow \lfloor \mathbf{M}_Q \rfloor_{q, \ell}$
- 9 $\ell \leftarrow \ell - \Delta^\rho / \Delta_s$
- 10 **end**

Algorithm 6: qDPMatrices

Data: q_0, T_f, \mathcal{G}
Result: $\mathbf{M}_Q, \mathbf{M}_p, \mathbf{M}_j$

- 1 $\alpha_{\max} \leftarrow \lfloor T_f / \Delta_s \rfloor$
- 2 $\mathbf{M}_Q \leftarrow [0] \in \{0, 1, \dots, Q(\delta)\}^{Q(\delta) \times (\alpha_{\max} + 1)}$ // $[\mathbf{M}_Q]_{q', \ell}$: best state q connected to q' at step ℓ
- 3 $\mathbf{M}_p \leftarrow [1] \in \{0, 1, \dots, D\}^{D \times (\alpha_{\max} + 1)}$ // $[\mathbf{M}_p]_{q', \ell}$: Perception connecting best q connected to q' at step ℓ
- 4 $\mathbf{M}_j \leftarrow [\infty] \in \mathbb{R}^{Q(\delta) \times (\alpha_{\max} + 1)}$ // $[\mathbf{M}_j]_{q, \ell}$: best cost from q_0 to q in ℓ steps
- 5 $[\mathbf{M}_j]_{q_0, 0} \leftarrow 0$
- 6 **for** $\ell \in \{1, \dots, \alpha_{\max}\}$ **do**
- 7 **for** $q \in \{1, \dots, Q(\delta)\}$ **do**
- 8 **for** $\rho \in \{1, \dots, D\}$ **do**
- 9 $q' \leftarrow$ state connected to q through edge with weight ρ
- 10 $\tau \leftarrow (\ell - 1)\Delta_s$
- 11 $\tau^+ \leftarrow \tau + \Delta^\rho$
- 12 $J \leftarrow \frac{1}{T_f} \left(\lambda_\alpha r^\rho + \int_\tau^{\min(\tau^+, T_f)} \mathbf{W}_d(t - \tau) dt + \int_\tau^{\min(\tau^+, T_f)} \text{tr}(\mathbf{A}_d(t - \tau) \hat{\mathbf{P}}_q \mathbf{A}_d(t - \tau)^\top) dt \right)$
 // $[\mathbf{M}_j]_{q, \ell} + J$ is the cost-to-arrive from q_0 to q' in $\ell + \rho$ steps
- 13 **if** $[\mathbf{M}_j]_{q, \ell} + J < [\mathbf{M}_j]_{q', \ell + \rho}$ **then**
- 14 $[\mathbf{M}_j]_{q', \ell + \rho} \leftarrow [\mathbf{M}_j]_{q, \ell} + J$
- 15 $[\mathbf{M}_Q]_{q', \ell + \rho} \leftarrow q$
- 16 $[\mathbf{M}_p]_{q', \ell + \rho} \leftarrow \rho$
- 17 **end**
- 18 **end**
- 19 **end**
- 20 **end**

Lemma 1. Let $\mathbf{M}_1, \mathbf{M}_2 \in \mathbb{R}^{n \times n}$ be positive definite matrices such that $\mathbf{M}_1 \preceq \mathbf{M}_2$. Then, the following inequalities are satisfied:

- (a) $\text{tr}(\mathbf{M}_1) \leq \text{tr}(\mathbf{M}_2)$.
- (b) $\|\mathbf{M}_1\|_F \leq \sqrt{n} \|\mathbf{M}_2\|_F$

Proof. First, note that $\mathbf{M}_2 - \mathbf{M}_1$ is positive semi-definite. Then, use the spectral Theorem [33, Theorem 2.5.6] to conclude that the eigenvalues $\{\lambda_i\}_{i=1}^n$ of $\mathbf{M}_2 - \mathbf{M}_1$ are all different non-negative

real numbers. Thus, $\text{tr}(\mathbf{M}_2) - \text{tr}(\mathbf{M}_1) = \text{tr}(\mathbf{M}_2 - \mathbf{M}_1) = \sum_{i=1}^n \lambda_i \geq 0$, where the relation between the trace and the sum of the eigenvalues was used in the last step [33, Page 50] completing the proof for a. Similarly, $\mathbf{M}_1, \mathbf{M}_2$ have non-negative eigenvalues too. Let $\lambda_1, \lambda_2 \in \mathbb{R}^n$ be vectors containing the eigenvalues of $\mathbf{M}_1, \mathbf{M}_2$ respectively. Thus, item (a) reads $\|\lambda_1\|_1 \leq \|\lambda_2\|_1$ using the notation of Proposition 5 in Appendix B. Proposition 5-a implies $\|\lambda_2\|_1 \leq \sqrt{n} \|\lambda_2\|_2$. Proposition 5-b implies $\|\lambda_1\|_2 \leq \|\lambda_1\|_1$. Thus, $\|\lambda_1\|_2 \leq \sqrt{n} \|\lambda_2\|_2$. Now, recall that $\|\lambda_1\|_2 = \sqrt{\sum_{i=1}^n \lambda_i(\mathbf{M}_1)^2} \equiv \|\mathbf{M}_1\|_F$ where $\lambda_i(\mathbf{M}_1)$ are the eigenvalues of \mathbf{M}_1 and similarly for \mathbf{M}_2 [33, Page 342]. Thus, item (b) follows.

Lemma 2. Let $\mathbf{M}_1, \mathbf{M}_2 \in \mathbb{R}^{n \times n}$ be symmetric matrices satisfying $\mathbf{M}_1 \preceq \mathbf{M}_2$. Then, $(\mathbf{M}_1 \otimes \mathbf{M}_1) \preceq (\mathbf{M}_2 \otimes \mathbf{M}_2)$.

Proof. First, recall from [33, Theorem 7.7.3-(a)] that $\mathbf{M}_1 \preceq \mathbf{M}_2$ if and only if $\rho(\mathbf{M}_2^{-1} \mathbf{M}_1) \leq 1$ where $\rho(\bullet)$ denotes the spectral radius [33, Definition 1.2.9]. Then, $1 \geq \rho(\mathbf{M}_2^{-1} \mathbf{M}_1)^2 = \rho((\mathbf{M}_2^{-1} \mathbf{M}_1) \otimes (\mathbf{M}_2^{-1} \mathbf{M}_1)) = \rho((\mathbf{M}_2 \otimes \mathbf{M}_2)^{-1} (\mathbf{M}_1 \otimes \mathbf{M}_1))$, which implies $(\mathbf{M}_1 \otimes \mathbf{M}_1) \preceq (\mathbf{M}_2 \otimes \mathbf{M}_2)$ using [33, Theorem 7.7.3-(a)].

Appendix C. Proofs**C.1. Proof of Theorem 1**

First, we show an following auxiliary result.

Lemma 3. Let the state Eq. (3) and the measurement model $\mathbf{z}[k] = \mathbf{C}\mathbf{x}[k] + \mathbf{v}[k]$ available at $t = \tau_k + \Delta^{p_k}$ with $\text{cov}\{\mathbf{v}[k]\} = \mathbf{R}^{p_k}$ and a fixed latency schedule p . Thus, the estimate at time $t \in [\tau_k, \tau_{k+1})$ of $\mathbf{x}(t)$ based on measurements $\{\mathbf{z}[0], \dots, \mathbf{z}[k]\}$ which minimize $\text{tr}(\hat{\mathbf{P}}(t))$, is given by $\mathbb{E}\{\mathbf{x}(t) | \mathbf{z}[0], \dots, \mathbf{z}[k-1]\}$ and satisfies (5).

Proof. First, consider estimates for $\mathbf{x}[k]$, given that $\mathbf{x}[k]$ evolves according to the discrete-time system (3b). Moreover, note that the measurement $\mathbf{z}[k]$ is available at $t = \tau_{k+1}$. Then, [34, Page 228 - Theorem 4.1] implies that the structure in (6), inherited from a Kalman filter with predictor, satisfies $\hat{\mathbf{x}}[k+1] \equiv \mathbb{E}\{\mathbf{x}[k+1] | \mathbf{z}[0], \dots, \mathbf{z}[k]\}$. For any other $t \in (\tau_k, \tau_{k+1})$ the measurement $\mathbf{z}[k]$ is not available. Thus, $\mathbb{E}\{\mathbf{x}(t) | \mathbf{z}[0], \dots, \mathbf{z}[k-1]\}$ can be computed using $\mathbb{E}\{\mathbf{x}[k] | \mathbf{z}[0], \dots, \mathbf{z}[k-1]\}$ through the same structure in (6) but applied to (3a) with $\mathbf{C} = 0$ resulting in (5). Now that $\hat{\mathbf{x}}(t) = \mathbb{E}\{\mathbf{x}(t) | \mathbf{z}[0], \dots, \mathbf{z}[k-1]\}$ has been established for any $t \in [\tau_k, \tau_{k+1})$, [34, Page 228 - Theorem 4.1] implies that $\mathbf{a}^\top \hat{\mathbf{P}}(t) \mathbf{a}$ is minimized for this estimate with arbitrary vector $\mathbf{a} \in \mathbb{R}^{n_x}$. This means that for any other estimation $\hat{\mathbf{x}}'(t)$ of $\mathbf{x}(t)$ with covariance $\hat{\mathbf{P}}'(t) := \text{cov}\{\mathbf{x}(t) - \hat{\mathbf{x}}'(t)\}$ we have $\mathbf{a}^\top \hat{\mathbf{P}}(t) \mathbf{a} \leq \mathbf{a}^\top \hat{\mathbf{P}}'(t) \mathbf{a}, \forall \mathbf{a} \in \mathbb{R}^{n_x}$. Hence, $\hat{\mathbf{P}}(t) \preceq \hat{\mathbf{P}}'(t)$ from which $\text{tr}(\hat{\mathbf{P}}(t)) \leq \text{tr}(\hat{\mathbf{P}}'(t))$ follows using Lemma 1 from Appendix B for any other estimation $\mathbf{x}'(t)$.

We are now ready to show Theorem 1. First, note that for fixed perception schedule p , the term $\sum_{k=0}^{\alpha} r^{p_k}$ in (2) is constant. Moreover, Lemma 3 implies that $\hat{\mathbf{x}}(t) = \mathbb{E}\{\mathbf{x}(t) | \mathbf{z}[0], \dots, \mathbf{z}[k-1]\}$ computed recursively using (5) and (6) leads to an optimal trajectory of the signal $\text{tr}(\hat{\mathbf{P}}(t))$. This means that for any other estimation $\hat{\mathbf{x}}'(t)$ of $\mathbf{x}(t)$ with $\hat{\mathbf{P}}'(t) := \text{cov}\{\mathbf{x}(t) - \hat{\mathbf{x}}'(t)\}$ one has $\text{tr}(\hat{\mathbf{P}}(t)) \leq \text{tr}(\hat{\mathbf{P}}'(t))$. Integrating both sides of the previous inequality leads to conclude that $\int_0^{T_f} \text{tr}(\hat{\mathbf{P}}(t)) dt \leq \int_0^{T_f} \text{tr}(\hat{\mathbf{P}}'(t)) dt$ which implies $\mathcal{J}(\hat{\mathbf{x}}, p) \leq \mathcal{J}(\hat{\mathbf{x}}', p)$. Hence, the estimation $\hat{\mathbf{x}}(t), \forall t \in [0, T_f]$ from (5) solves Problem 1 provided that p is the optimal schedule as well.

C.2. Proof of Proposition 2

To show correctness, first note that a direct application of the dynamic programming principle [36, Page 23] leads to conclude that p as obtained from $\text{dynProg}(0, P[0], T_f)$ is optimal for the cost (2) whenever the estimations $\{\hat{\mathbf{x}}(t) : t \in [0, T_f]\}$ are of the form (5) as required in lines 4 and 6 of dynProg . In addition, Theorem 1 implies that the optimal estimations are computed from (5), from which optimality of both $\hat{\mathbf{x}}(t)$ and p using dynProg follows, solving Problem 1. For the complexity, note that the largest amount of recursive calls in Algorithm 2 is obtained when $p_k = i$ at line 3 of dynProg with $i = \text{argmin}\{\Delta^i\}_{j=1}^D$. In this case, $\lceil T_f / \min\{\Delta^1, \dots, \Delta^D\} \rceil$ recursive calls to dynProg are required to cover the whole window $[0, T_f]$, i.e. for $\tau^+ \geq T_f$ in line 7 of dynProg . Hence, due to line 3 in dynProg , the total number of recursive calls is at most $D \lceil T_f / \min\{\Delta^1, \dots, \Delta^D\} \rceil$.

C.3. Proof of Proposition 3

First we show that Algorithm 3 finishes. Note that for any initial condition in a compact set \mathcal{B}_0 , the result in [47, Lemma 6.1] ensures that the covariance $\hat{\mathbf{P}}[k]$ a the filter of the form (6) is uniformly bounded from above. This means that there exists a compact set $\mathcal{B} \subset \mathbb{R}^{n_x \times n_x}$ such that $\hat{\mathbf{P}}[k] \in \mathcal{B}, \forall k \geq 0$ and $\hat{\mathbf{P}}_0 \in \mathcal{B}_0$. Existence of such compact set in the quantized setting is ensured as well. The reason is that quantization only adds a disturbance term $\|\mathbf{II}_Q[k]\| \leq \delta$ to (6) which can be absorbed into the covariance $W(\Delta^{p_k})$. Now, note that since quantization patches obtained in Algorithm 3 are non overlapping and comply $\sup_{\hat{\mathbf{P}}[k], \hat{\mathbf{P}}[k'] \in \mathcal{B}_{Q(\delta)}} \|\hat{\mathbf{P}}[k] - \hat{\mathbf{P}}[k']\| = \delta$ thus, there exists a maximum finite number $\bar{Q}(\delta)$ of regions of this kind covering \mathcal{B} due to compactness. Hence, Algorithm 3 finishes in at most $\bar{Q}(\delta)$ steps.

For the rest of the proposition, the proof outline is described in the following. Note that obtaining an explicit bound for $\hat{\mathbf{P}}[k]$ in (6) is not trivial due to the nonlinearity of the update equation for the covariance. In addition, the bound in [47, Lemma 6.1] is very overestimated. Hence, we follow a similar idea as in [47] which is to study an artifact filter whose bound can be obtained explicitly and use it as a proof tool, appealing to the optimality of PLATE to conclude that a covariance bound of the artifact filter is a bound for PLATE as well. In this case, instead of the structure in (6) where the gain $\mathbf{L}[k]$ depends on $\hat{\mathbf{P}}[k]$, we use a gain $\mathbf{L}^{p_k} \in \{\mathbf{L}^1, \dots, \mathbf{L}^D\}$ so that the artifact filter takes the form:

$$\hat{\mathbf{x}}_s[k+1] = \mathbf{A}_d(\Delta^{p_k})\hat{\mathbf{x}}_s[k] + \mathbf{L}^{p_k}(\mathbf{z}[k] - \mathbf{C}\hat{\mathbf{x}}_s[k]) \quad (\text{C.1})$$

with $\hat{\mathbf{x}}_s[0] = \mathbf{x}_0$. Hence, we require to design the gains \mathbf{L}^{p_k} to make the filter asymptotically stable, making it feasible for the bound to exist. This cannot be done by designing each \mathbf{L}^{p_k} by separate. The reason is that, the error $\tilde{\mathbf{x}}_s[k] = \mathbf{x}[k] - \hat{\mathbf{x}}_s[k]$ for (C.1) satisfies

$$\tilde{\mathbf{x}}_s[k+1] = \mathbf{A}^{p_k}\tilde{\mathbf{x}}_s[k] + \mathbf{L}^{p_k}\mathbf{v}[k] + \mathbf{w}_d[k] \quad (\text{C.2})$$

with $\mathbf{A}^{p_k} = \mathbf{A}_d(\Delta^{p_k}) + \mathbf{L}^{p_k}\mathbf{C}$. Hence, (C.2) is a switched system which switches between system matrices \mathbf{A}^{p_k} with the schedule p_k as switching signal. In the following auxiliary technical lemmas, we aim to ensure asymptotic stability of the filter, by designing \mathbf{L}^{p_k} through the method of the common Lyapunov function [48, Page 22].

Lemma 4. Consider the filter (C.1) for some perception schedule p . Then, $\mathbf{q}_s[k] := \text{vec}(\hat{\mathbf{P}}_s[k])$ with $\hat{\mathbf{P}}_s[k] := \text{cov}\{\mathbf{x}[k] - \mathbf{x}_s[k]\}$ satisfies:

$$\mathbf{q}_s[k+1] = (\mathbf{A}^{p_k} \otimes \mathbf{A}^{p_k})\mathbf{q}_s[k] + \boldsymbol{\omega}[k], \mathbf{q}_s[0] = \text{vec}(\mathbf{P}_0) \quad (\text{C.3})$$

where $\boldsymbol{\omega}[k] = \text{vec}((\mathbf{L}^{p_k}\mathbf{R}^{p_k}(\mathbf{L}^{p_k})^\top + \mathbf{W}_d(\Delta^{p_k}))$.

Proof. First, compute $\hat{\mathbf{P}}_s[k+1]$ by applying $\text{cov}(\bullet)$ to both sides of (C.2) as:

$$\begin{aligned} \hat{\mathbf{P}}_s[k+1] &= (\mathbf{A}^{p_k})\hat{\mathbf{P}}_s[k](\mathbf{A}^{p_k})^\top \\ &\quad + (\mathbf{L}^{p_k})\mathbf{R}^{p_k}(\mathbf{L}^{p_k})^\top + \mathbf{W}_d(\Delta^{p_k}) \end{aligned}$$

Then, apply $\text{vec}(\bullet)$ to both sides of the previous equation as well as the identity $\text{vec}((\mathbf{A}^{p_k})\hat{\mathbf{P}}_s[k](\mathbf{A}^{p_k})^\top) = (\mathbf{A}^{p_k} \otimes \mathbf{A}^{p_k})\text{vec}(\hat{\mathbf{P}}_s[k])$ to obtain (C.3). Finally, $\hat{\mathbf{P}}[0] = \text{cov}\{\mathbf{x}[0] - \mathbf{x}_0\} = \mathbf{P}_0$.

Lemma 5. Consider $\boldsymbol{\Omega}, \gamma, \mathbf{Y}^i, i \in \{1, \dots, D\}$ satisfy (8). Henceforth, the following nonlinear matrix inequality is satisfied:

$$((\mathbf{A}^{p_k})^\top \boldsymbol{\Omega}(\mathbf{A}^{p_k}) \otimes (\mathbf{A}^{p_k})^\top \boldsymbol{\Omega}(\mathbf{A}^{p_k})) \leq \gamma^2(\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}) \quad (\text{C.4})$$

with $\mathbf{A}^i = \mathbf{A}_d(\Delta^i) - \mathbf{L}^i\mathbf{C}$ and $\mathbf{L}^i = \boldsymbol{\Omega}^{-1}\mathbf{Y}^i, i \in \{1, \dots, D\}$.

Proof. Equivalence between (8) and $(\mathbf{A}^{p_k})^\top \boldsymbol{\Omega}(\mathbf{A}^{p_k}) \leq \gamma \boldsymbol{\Omega}$ follows from the well known relationship of the Schur complement similarly as in [49]. Now, apply Lemma 2 with $\mathbf{M}_1 = (\mathbf{A}^{p_k})^\top \boldsymbol{\Omega}(\mathbf{A}^{p_k})$ and $\mathbf{M}_2 = \gamma \boldsymbol{\Omega}$ to obtain (C.4).

Lemma 6. Consider the assumptions in Lemmas 4 and 5. Moreover, let $V(\mathbf{q}_s[k]) := \sqrt{\mathbf{q}_s[k]^\top (\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}) \mathbf{q}_s[k]}$. Then, the following inequality is satisfied:

$$V(\mathbf{q}_s[k+1]) \leq \gamma V(\mathbf{q}_s[k]) + \lambda_{\max}(\boldsymbol{\Omega})\bar{G} \quad (\text{C.5})$$

where $\lambda_{\max}(\boldsymbol{\Omega})$ is the maximum eigenvalue of $\boldsymbol{\Omega}$ and

$$\bar{G} := \max_{i \in \{1, \dots, D\}} \|(\mathbf{L}^{p_k})\mathbf{R}^{p_k}(\mathbf{L}^{p_k})^\top + \mathbf{W}_d(\Delta^{p_k})\|_F.$$

Proof. First, note that since $\boldsymbol{\Omega}$ is positive definite, henceforth $\|\bullet\|_{\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}} := \sqrt{(\bullet)^\top (\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}) (\bullet)}$ is a norm [33, Page 321] and $V(\mathbf{q}_s[k]) \equiv \|\mathbf{q}_s[k]\|_{\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}}$. Now compute $V(\mathbf{q}_s[k+1])$ from (C.3) as:

$$\begin{aligned} V(\mathbf{q}_s[k+1]) &= \|(\mathbf{A}^{p_k} \otimes \mathbf{A}^{p_k})\mathbf{q}_s[k] + \boldsymbol{\omega}[k]\|_{\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}} \\ &\leq \|(\mathbf{A}^{p_k} \otimes \mathbf{A}^{p_k})\mathbf{q}_s[k]\|_{\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}} + \|\boldsymbol{\omega}[k]\|_{\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}} \end{aligned}$$

using the triangle inequality [33, Definition 5.1.1-(3)]. Furthermore, use Lemma 5 to obtain:

$$\begin{aligned} &\|(\mathbf{A}^{p_k} \otimes \mathbf{A}^{p_k})\mathbf{q}_s[k]\|_{\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}} \\ &= \sqrt{\mathbf{q}_s[k]^\top ((\mathbf{A}^{p_k})^\top \boldsymbol{\Omega}(\mathbf{A}^{p_k}) \otimes (\mathbf{A}^{p_k})^\top \boldsymbol{\Omega}(\mathbf{A}^{p_k})) \mathbf{q}_s[k]} \\ &\leq \sqrt{\gamma^2 \mathbf{q}_s[k]^\top (\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}) \mathbf{q}_s[k]} \\ &= \gamma \|\mathbf{q}_s[k]\|_{\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}} \equiv \gamma V(\mathbf{q}_s[k]) \end{aligned} \quad (\text{C.6})$$

Moreover, note that $\boldsymbol{\omega}[k]^\top (\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}) \boldsymbol{\omega}[k] \leq \lambda_{\max}(\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}) \|\boldsymbol{\omega}[k]\|^2$ by the Rayleigh inequality [33, Theorem 4.2.2]. Furthermore, note that

$$\|\boldsymbol{\omega}[k]\| = \|(\mathbf{L}^{p_k})\mathbf{R}^{p_k}(\mathbf{L}^{p_k})^\top + \mathbf{W}_d(\Delta^{p_k})\|_F.$$

Thus, $\|\boldsymbol{\omega}[k]\|_{\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}} \leq \lambda_{\max}(\boldsymbol{\Omega})\bar{G}$. Then, (C.5) is the combination of the previous results.

We are now ready to show the rest of Proposition 3. First, consider a scalar signal $v[k] \in \mathbb{R}$ satisfying:

$$v[k+1] = \gamma v[k] + \lambda_{\max}(\boldsymbol{\Omega})\bar{G},$$

with $\lambda_{\max}(\boldsymbol{\Omega})\bar{G}$ as in Lemma 6, $\gamma \in (0, 1)$ and

$v[0] = \sqrt{\text{vec}(\mathbf{P}_0)^\top (\boldsymbol{\Omega} \otimes \boldsymbol{\Omega}) \text{vec}(\mathbf{P}_0)}$. It can be verified that the solution to the previous linear difference equation satisfies:

$$\begin{aligned} v[k] &= \gamma^k v[0] + \lambda_{\max}(\boldsymbol{\Omega})\bar{G} \left(\frac{1 - \gamma^k}{1 - \gamma} \right) \\ &\leq v[0] + \frac{\lambda_{\max}(\boldsymbol{\Omega})\bar{G}}{1 - \gamma}, \forall k \geq 0 \end{aligned}$$

In addition, note that

$$v[0] \leq \sqrt{\lambda_{\max}(\Omega \otimes \Omega) \text{vec}(\mathbf{P}_0)^\top \text{vec}(\mathbf{P}_0)} = \lambda_{\max}(\Omega) \|\mathbf{P}_0\|_F$$

by means of the Rayleigh inequality [33, Theorem 4.2.2] and the identity $\|\text{vec}(\mathbf{P}_0)\| \equiv \|\mathbf{P}_0\|_F$. Combine the previous results to conclude that $v[k] \leq B'_s, \forall k \geq 0$ with $B'_s = \lambda_{\max}(\Omega) \left(\|\mathbf{P}_0\|_F + \frac{\bar{c}}{1-\gamma} \right)$. Recall that $V(\mathbf{q}_s[k])$ complies (C.5) by Lemma 6. Use the comparison lemma in [50, Lemma 13] together with $V(\mathbf{q}_s[0]) = v[0]$ to conclude that $V(\mathbf{q}_s[k]) \leq v[k], \forall k \geq 0$ which leads directly $V(\mathbf{q}_s[k]) \leq B'_s$. Now, Rayleigh inequality is used again to conclude that $\lambda_{\min}(\Omega) \|\mathbf{q}_s[k]\| \leq V(\mathbf{q}_s[k]) \leq B'_s$ equivalently $\|\mathbf{q}_s[k]\| = \|\hat{\mathbf{P}}_s[k]\|_F \leq B'_s / \lambda_{\min}(\Omega) \equiv B_s / \sqrt{n_x}$ with B_s defined in (7). Use the same arguments as in the proof of Lemma 3 to conclude that $\hat{\mathbf{P}}[k] \preceq \hat{\mathbf{P}}_s[k]$ where $\hat{\mathbf{P}}[k]$ is obtained from PLATE in (6). Hence, use the same arguments as in the proof of Lemma 1-b in Appendix B to obtain $\|\hat{\mathbf{P}}[k]\|_F \leq \sqrt{n_x} \|\hat{\mathbf{P}}_s[k]\|_F \leq B_s$.

C.4. Proof of Proposition 4

Similar to the proof of Proposition 2, optimality of the solution of Algorithm 5 comes from the dynamic programming principle [36, Page 23]. As for the complexity, note that combining lines 6, 7 and 9 of Algorithm 6 results in $\alpha_{\max} Q(\delta) D$ iterations needed to compute $\mathbf{M}_p, \mathbf{M}_j, \mathbf{M}_a$. In addition, the number of steps in lines 2 and 5 in Algorithm 5 are $Q(\delta)$ and α_{\max} respectively. Hence, the asymptotic complexity of Algorithm 5 is only given by the term $\alpha_{\max} Q(\delta) D$.

C.5. Proof of Theorem 2

In this section, we denote with \mathbf{F}^{p_k} the function that computes $\hat{\mathbf{P}}[k+1] = \mathbf{F}^{p_k}(\hat{\mathbf{P}}[k])$ according to (6). Moreover, in order to show Theorem 2, we provide an auxiliary lemma.

Lemma 7.

Let a given $K \in \mathbb{N}$ and consider the systems:

$$\hat{\mathbf{P}}[k+1] = \mathbf{F}^{p_k}(\hat{\mathbf{P}}[k]), \quad \hat{\mathbf{P}}[0] = \mathbf{P}_0$$

$$\hat{\mathbf{P}}'[k+1] = \mathbf{F}^{p_k}(\hat{\mathbf{P}}'[k]) + \mathbf{II}_Q[k], \quad \hat{\mathbf{P}}'[0] = \mathbf{P}'_0$$

with $\|\hat{\mathbf{P}}[0] - \hat{\mathbf{P}}'[0]\|_F \leq \delta$ and $\|\mathbf{II}_Q[k]\|_F \leq \delta, \forall k \in \{1, \dots, K\}$. Therefore, for any $\varepsilon' > 0$ there exists $\delta > 0$ such that $\|\hat{\mathbf{P}}[k] - \hat{\mathbf{P}}'[k]\|_F \leq \varepsilon', \forall k \in \{1, \dots, K\}$.

Proof. First, note that for any $\varepsilon_0 > 0$, there exists $\delta > 0$ such that $\|\hat{\mathbf{P}}[1] - \hat{\mathbf{P}}'[1]\|_F = \|\mathbf{F}^{p_0}(\hat{\mathbf{P}}[0]) - \mathbf{F}^{p_0}(\hat{\mathbf{P}}'[0]) - \mathbf{II}_Q[0]\|_F \leq \|\mathbf{F}^{p_0}(\hat{\mathbf{P}}[0]) - \mathbf{F}^{p_0}(\hat{\mathbf{P}}'[0])\|_F + \delta \leq \varepsilon_0$ due to continuity of $\mathbf{F}^{p_0}(\bullet)$. The same reasoning applies for the remaining $K-1$ steps, making $\hat{\mathbf{P}}[k]$ arbitrarily close to $\hat{\mathbf{P}}'[k]$ by choosing $\delta > 0$ sufficiently small.

Using the previous result, the proof of Theorem 2 follows. First note that the initial condition \mathbf{P}_0 for the original problem and the quantized version $\hat{\mathbf{P}}_{q_0} = \mathcal{Q}(\mathbf{P}_0)$ comply $\|\mathbf{P}_0 - \hat{\mathbf{P}}_{q_0}\|_F \leq \delta$. Note that if $\hat{\mathbf{P}}[k] \in \{\hat{\mathbf{P}}_1, \dots, \hat{\mathbf{P}}_{Q(\delta)}\}$, then $\|\mathbf{F}^{p_k}(\hat{\mathbf{P}}[k]) - \mathcal{Q}(\mathbf{F}^{p_k}(\hat{\mathbf{P}}[k]))\|_F \leq \delta$. Thus, in the quantized setting, $\hat{\mathbf{P}}'[k]$ evolves according to $\hat{\mathbf{P}}'[k+1] = \mathbf{F}^{p_k}(\hat{\mathbf{P}}'[k]) + \mathbf{II}_Q[k]$ where $\mathbf{II}_Q[k]$ is the quantization noise complying $\|\mathbf{II}_Q[k]\|_F \leq \delta$ and initial condition $\hat{\mathbf{P}}'[0] = \hat{\mathbf{P}}_{q_0}$. Thus, Lemma 7 implies that for any $\varepsilon' > 0$ there is sufficiently small $\delta > 0$ such that true optimal trajectory for the covariance $\hat{\mathbf{P}}[k]$ and the quantized one $\hat{\mathbf{P}}'[k]$ comply $\|\hat{\mathbf{P}}[k] - \hat{\mathbf{P}}'[k]\|_F \leq \varepsilon'$ equivalently $\hat{\mathbf{P}}[k] - \hat{\mathbf{P}}'[k] = \hat{\mathbf{P}}[k]$ for some $\hat{\mathbf{P}}[k]$ with $\|\hat{\mathbf{P}}[k]\|_F \leq \varepsilon'$.

Now, with $\tau_{k+1}^+ := \min(\tau_{k+1}, T_f)$:

$$\begin{aligned} |\mathcal{J} - \mathcal{J}_Q(\delta)| &\leq \frac{1}{T_f} \sum_{k=0}^{\text{len}(p)} \int_{\tau_k}^{\tau_{k+1}^+} \left| \text{tr}(\mathbf{A}_d(t - \tau_k) \tilde{\mathbf{P}}[k] \mathbf{A}_d(t - \tau_k)^\top) \right| dt \\ &\leq \frac{1}{T_f} \sum_{k=0}^{\text{len}(p)} (\tau_{k+1}^+ - \tau_k) \sup_{\tau_k \leq \tau \leq \tau_{k+1}^+} \left| \text{tr}(\mathbf{A}_d(\tau) \tilde{\mathbf{P}}[k] \mathbf{A}_d(\tau)^\top) \right| \\ &\leq \max_{0 \leq k \leq \text{len}(p)} \sup_{\tau_k \leq \tau \leq \tau_{k+1}^+} \left| \text{tr}(\mathbf{A}_d(\tau) \tilde{\mathbf{P}}[k] \mathbf{A}_d(\tau)^\top) \right| \end{aligned}$$

However, since $\|\tilde{\mathbf{P}}[k]\|_F \leq \varepsilon'$, one can choose $\delta > 0$ sufficiently small to make $\varepsilon' > 0$ and as a consequence $\left| \text{tr}(\mathbf{A}_d(\tau) \tilde{\mathbf{P}}[k] \mathbf{A}_d(\tau)^\top) \right|$ to be arbitrarily small for any $\tau \in [\tau_k, \tau_{k+1}], k \in \{1, \dots, \text{len}(p)\}$. Thus, $|\mathcal{J} - \mathcal{J}_Q(\delta)| \leq \varepsilon$ for some $\delta > 0$.

References

- [1] Luo W, Xing J, Milan A, Zhang X, Liu W, Kim T-K. Multiple object tracking: A literature review. *Artificial Intelligence* 2021;293:103448.
- [2] Hu W, Tan T, Wang L, Maybank S. A survey on visual surveillance of object motion and behaviors. *IEEE Trans Syst Man Cybern Part C Appl Rev* 2004;34(3):334–52.
- [3] Madrigal F, Hayet J-B. Motion priors based on goals hierarchies in pedestrian tracking applications. *Mach Vis Appl* 2017;28:341–59.
- [4] Mitra S, Acharya T. Gesture recognition: A survey. *IEEE Trans Syst Man Cybern Part C Appl Rev* 2007;37:311–24.
- [5] Datondji SRE, Dupuis Y, Subirats P, Vasseur P. A survey of vision-based traffic monitoring of road intersections. *IEEE Trans Intell Transp Syst* 2016;17:2681–98.
- [6] Chen P, Dang Y, Liang R, Zhu W, He X. Real-time object tracking on a drone with multi-inertial sensing data. *IEEE Trans Intell Transp Syst* 2018;19(1):131–9.
- [7] Khan A, Rinner B, Cavallaro A. Cooperative robots to observe moving targets: Review. *IEEE Trans Cybern* 2018;48(1):187–98.
- [8] Bescos B, Campos C, Tardós JD, Neira J. Dynaslam ii: Tightly-coupled multi-object tracking and slam. *IEEE Rob Autom Lett* 2021;6(3):5191–8.
- [9] Yilmaz A, Javed O, Shah M. Object tracking: A survey. *ACM Comput Surv* 2006;38(4):13–es.
- [10] Redmon J, Divvala SK, Girshick RB, Farhadi A. You only look once: Unified, real-time object detection. In: *PROC CVPR IEEE*. 2016, p. 779–88.
- [11] Beery S, Wu G, Rathod V, Votel R, Huang J. Context r-cnn: Long term temporal context for per-camera object detection. In: *PROC CVPR IEEE*. 2020, p. 13072–82.
- [12] Falanga D, Kim S, Scaramuzza D. How fast is too fast? the role of perception latency in high-speed sense and avoid. *IEEE Rob Autom Lett* 2019;4(2):1884–91.
- [13] Huang J, Rathod V, Sun C, Zhu M, Korattikara A, Fathi A, et al. Speed/accuracy trade-offs for modern convolutional object detectors. In: *PROC CVPR IEEE*. 2017, p. 3296–7.
- [14] Narykov AS, Krasnov OA, Yarovsky A. Algorithm for resource management of multiple phased array radars for target tracking. In: *Proceedings of the 16th international conference on information fusion*. 2013, p. 1258–64.
- [15] Gatsis K, Hassani H, Pappas GJ. Latency-reliability tradeoffs for state estimation. *IEEE Trans Automat Control* 2021;66(3):1009–23. <http://dx.doi.org/10.1109/TAC.2020.2992563>.
- [16] Fan H, Ling H. Parallel tracking and verifying: A framework for real-time and high accuracy visual tracking. In: *IEEE international conference on computer vision*. 2017, p. 5487–95.
- [17] Howard AG, Zhu M, Chen B, Kalenichenko D, Wang W, Weyand T, et al. Mobilenets: Efficient convolutional neural networks for mobile vision applications. 2017, arXiv:abs/1704.04861.
- [18] Casares M, Velipalasar S. Adaptive methodologies for energy-efficient object detection and tracking with battery-powered embedded smart cameras. *IEEE Trans Circ Syst Video Technol* 2011;21(10):1438–52.
- [19] Guan M, Wen C, Shan M, Ng C-L, Zou Y. Real-time event-triggered object tracking in the presence of model drift and occlusion. *IEEE Trans Ind Electron* 2019;66(3):2054–65.
- [20] Hu H, Dey D, Hebert M, Bagnell J. Learning anytime predictions in neural networks via adaptive loss balancing. In: *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 2019, p. 3812–21.
- [21] Yao S, Hao Y, Zhao Y, Shao H, Liu D, Liu S, et al. Scheduling real-time deep learning services as imprecise computations. In: *IEEE 26th international conference on embedded and real-time computing systems and applications*. 2020.

- [22] Aldana-López R, Aragués R, Sagiés C. Attention vs. precision: Latency scheduling for uncertainty resilient control systems. In: 2020 59th IEEE conference on decision and control. 2020, p. 5697–702.
- [23] Luo H, Xie W, Wang X, Zeng W. Detect or track: Towards cost-effective video object detection/tracking. In: Proceedings of the AAAI conference on artificial intelligence, Vol. 33. 2019, p. 8803–10.
- [24] Wu X, Zhang K, Sun C. Optimal scheduling of multiple sensors in continuous time. *ISA Trans* 2014;53(3):793–801.
- [25] Sid M. Sensor scheduling strategies for fault isolation in networked control system. *ISA Trans* 2015;54:92–100.
- [26] Zhao L, Zhang W, Hu J, Abate A, Tomlin CJ. On the optimal solutions of the infinite-horizon linear sensor scheduling problem. *IEEE Trans Automat Control* 2014;59(10):2825–30.
- [27] Huber MF. Optimal pruning for multi-step sensor scheduling. *IEEE Trans Automat Control* 2012;57(5):1338–43.
- [28] Asghar AB, Jawaid ST, Smith SL. A complete greedy algorithm for infinite-horizon sensor scheduling. *Automatica* 2017;81:335–41.
- [29] Orihuela L, Barreiro A, Gómez-Estern F, Rubio FR. Periodicity of kalman-based scheduled filters. *Automatica* 2014;50(10):2672–6.
- [30] Li Y, Chen C, Zhu S, Guan X. Sensor scheduling for relay-assisted wireless control systems with limited power resources. *ISA Trans* 2019;88: 246–57.
- [31] Arai S, Iwatani Y, Hashimoto K. Fast sensor scheduling for spatially distributed sensors. *IEEE Trans Automat Control* 2011;56: (8):1900–5.
- [32] Qi Y, Cheng P, Chen J. Optimal sensor data scheduling for remote estimation over a time-varying channel. *IEEE Trans Automat Control* 2017;62(9):4611–7.
- [33] Horn RA, Johnson CR. Matrix analysis. 2nd ed.. USA: Cambridge University Press; 2012.
- [34] Åström K. Introduction to stochastic control theory, mathematics in science and engineering. Academic Press; 1970.
- [35] Soderstrom T. Discrete-time stochastic systems: estimation and control. 2nd ed.. Berlin, Heidelberg: Springer-Verlag; 2002.
- [36] Bertsekas DP. Dynamic programming and optimal control. 2nd ed.. Athena Scientific; 2000.
- [37] Song Y-M, Yoon K, Yoon Y-C, Yow KC, Jeon M. Online multi-object tracking with gmphd filter and occlusion group management. *IEEE Access* 2019;7:165103–21.
- [38] Dong X, Shen J, Yu D, Wang W, Liu J, Huang H. Occlusion-aware real-time object tracking. *IEEE Trans Multimed* 2017;19(4):763–71.
- [39] Wang L, Zhang L, Yi Z. Trajectory predictor by using recurrent neural networks in visual tracking. *IEEE Trans Cybern* 2017;47(10):3172–83.
- [40] Loquercio A, Segu M, Scaramuzza D. A general framework for uncertainty estimation in deep learning. *IEEE Rob Autom Lett* 2020;5(2):3153–60.
- [41] Fokin LA, Shchipitsyn AG. Innovation-based adaptive kalman filter derivation. In: 2009 international siberian conference on control and communications. 2009, p. 318–23.
- [42] Milan A, Leal-Taixé L, Reid I, Roth S, Schindler K. MOT16: A benchmark for multi-object tracking. 2016, arXiv:1603.00831.
- [43] Wu Y, Kirillov A, Massa F, Lo W-Y, Girshick R. Detectron2. 2019, <https://github.com/facebookresearch/detectron2>.
- [44] Zhou K, Yang Y, Cavallaro A, Xiang T. Omni-scale feature learning for person re-identification. In: 2019 IEEE/CVF international conference on computer vision. 2019, p. 3701–11.
- [45] Bertinetto L, Valmadre J, Henriques JF, Vedaldi A, Torr PHS. Fully-convolutional siamese networks for object tracking. In: Hua G, Jégou H, editors. Computer vision – ECCV 2016 workshops. Cham: Springer International Publishing; 2016, p. 850–65.
- [46] Hardy G, Littlewood J, Pólya G. Inequalities. Cambridge Mathematical Library, Cambridge University Press; 1988.
- [47] Moore JB, Anderson BDO. Coping with singular transition matrices in estimation and control stability theory. *Internat J Control* 1980;31(3):571–86.
- [48] Liberzon D. Switching in systems and control. Systems & Control: Foundations & Applications, Birkhäuser Boston; 2003.
- [49] Alessandri A, Coletta P. Switching observers for continuous-time and discrete-time linear systems. In: Proceedings of the 2001 american control conference. (Cat. No. 01CH37148), Vol. 3. 2001, p. 2516–21.
- [50] Jiménez-Rodríguez E, Aldana-López R, Sánchez-Torres JD, Gómez-Gutiérrez D, Loukianov AG. Consistent discretization of a class of predefined-time stable systems. *IFAC-PapersOnLine* 2020;53(2):628–33, 21st IFAC World Congress.