Concealing CAN Message Sequences to Prevent Schedule-based Bus-off Attacks

Sunandan Adhikary *, Ipsita Koley *, Arkaprava Sain*, Soumyadeep das ‡, Shuvam Saha §, Soumyajit Dey *

* IIT Kharagpur, India ,‡NIT Durgapur, India ,§Jadavpur University, India

Abstract—This work focuses on eliminating timing-side channels in real-time safety-critical cyber-physical network protocols like Controller Area Networks (CAN). Automotive Electronic Control Units (ECUs) implement predictable scheduling decisions based on task level response time estimation. Such levels of determinism exposes timing information about task executions and therefore corresponding message transmissions via the network buses (that connect the ECUs and actuators). With proper analysis, such timing side channels can be utilized to launch several schedule-based attacks that can lead to eventual denial-of-service or man-in-the-middle-type attacks. To eliminate this determinism, we propose a novel schedule obfuscation strategy by *skipping* certain control task executions and related data transmissions along with *random shifting* of the victim task instance. While doing this, our strategy contemplates the performance of the control task as well by bounding the number of control execution skips. We analytically demonstrate how the *attack success probability (ASP)* is reduced under this proposed *attack-aware skipping and randomization*. We also demonstrate the efficacy and real-time applicability of our *attack-aware schedule obfuscation strategy Hide-n-Seek* by applying it to synthesized automotive task sets in a real-time Hardware-in-loop (HIL) setup.

Index Terms—Timing Attacks, Control Execution Skips , CAN Security, Schedule Randomization

I. INTRODUCTION

The development cycle of real-time cyber-physical systems (CPSs), like automotive, generally considers safety and time-liness as their first-class design criteria. This along with the limited availability of computation/communication resources often lead to adoption of predictable static/dynamic task scheduling algorithms. As a result, we have deterministic task execution schedules (e.g. Autosar classic platforms) involving multiple control tasks for several CPS closed loops running in shared real-time embedded platforms. Achieving such pre-dictable timing and functionality while considering the *worst case* behaviour is also desired by the system designers in order to ensure the safety of the closed-loop in the presence of delay, jitters and interferences in shared processors and communication mediums. It may be noted that platform-level deterministic task execution leads to regular patterns of messages in the associated communication medium that the processing units use for sensing, actuation as well as inter-ECU data sharing. This predictability in task and message schedules open up timing side-channels [1] in processor and communication channels.

In modern automotive systems, modular design of the application layers, runtime environments, and communication stack designed following the AUTOSAR standards [2], highly restricts on-board vulnerabilities from exposure. Therefore, the lightweight in-vehicle communication protocols have become ideal targets of a *schedule-based attacker*. The most preferred target for such *man-in-the-middle* type attacks is the Control Area Network (CAN) [3], the de-facto intra-vehicular communication protocol that transmits most of the *real-time safety-critical messages* among the electronic control units (ECUs). Typically, tasks in vehicle ECUs follow a static schedule whenever arrival of new tasks does not occur. The related messages produced by the control tasks are priority-wise communicated via CAN bus following a static schedule. Since CAN protocol lacks any authentication and confidentiality scheme, it becomes an ideal target for an attacker. The deterministic timing behaviour of the safety-critical CAN messages is utilized to launch schedule-based attacks that can disconnect an ECU temporarily by sending it to *bus-off* mode (known as *bus-off attack*) [1], [4]. In a bus-off attack, the attacker is essentially a piece of code implemented on a compromised ECU. It analyzes the CAN traffic, identifies the periodicity and arrival times of a targeted message, fabricates an attack message with the same identifier as the targeted one. The attacker injects the attack message into CAN synchronously with the target message at the same time. Repeated synchronous transmissions cause the transmission error count of the victim message's source ECU to increase and gradually send the victim ECU to the bus-off mode. The bus-off attack is one of the most lethal attacks on CAN protocol proposed in the last decade because the attacker remains stealthy during the attack process as it uses the same identifier and timing information as the victim message. With the victim ECU in bus-off mode, the attacker can now transmit false data messages mimicking the victim and causing further damage to the system.

A significant number of solutions have been proposed to counter bus-off attacks. One line of work uses physical properties, like clock skew, and voltage profile, of the ECUs to identify the source of all the messages beforehand [5]–[7]. At run-time, if a message exhibits different values of such properties, it is flagged as coming from a compromised ECU. The primary limitation of these approaches is that they take significant time to detect bus-off attack attempts. On the other hand, once the synchronization with the victim message is achieved, the entire bus-off attack takes very less time to accomplish. Another line of intrusion detection systems proposed to encounter bus-off attacks is the use of trusted modules to provide authentication to the transmitted messages

[8]–[10]. However, these solutions incur additional installation costs. Moreover, if the trusted module is compromised, the entire intrusion detection system fails. In this work, we propose a lightweight attack-aware schedule randomization strategy *Hide-n-Seek* to obfuscate (*hide*) the CAN schedule from timing leakages as well as detect (*seek*) the bus-off attack attempts. *Schedule randomization* is largely suggested as a helpful strategy to defend against processor-level timing attacks [11], [12]. Randomized schedules can obfuscate the timing information about the schedule and theoretically promise higher entropy. However, it has also been shown by some researchers that the increment in the entropy of schedule randomization cannot ensure reduced *attack success probability (ASP)*. Rather randomizing without any insight into the attacker's strategy might increase the penetrability of the system [13].

In this work, we propose an attack-aware schedule obfuscation strategy (*AAwSOS*) that instruments an ECU schedule that hides the timing information of safety-critical messages, transmitted through the CAN bus, from an attacker who is attempting a schedule-based attack (SBA), like bus-off attack. The proposed *AAwSOS* module in each ECU e will first monitor the traffic for a certain number of CAN hyper-periods. Next, it will find out which instances of the safety-critical control tasks, transmitted from e, are most vulnerable to bus-off attack w.r.t. this analysis window. *AAwSOS* then either **(a)** skips those vulnerable instances of the safety-critical control tasks, or **(b)** skips one or more instances of some higher priority control tasks in e. In general, control theorists design mathematical control laws while keeping environmental and platform delays in mind. This generally implies the control loop will work with the desired performance guarantee up to a bounded number of delay-induced control task execution skips [14]. In this work, we propose to intentionally inject such skips with security as a goal in mind. Switching between many such control skipping sequences makes information leakage about the execution decision of task instances less probable as the positions of job skips for the same task keep changing. Consequently, the faulty schedule analysis exposes the attacker. However, skipping an arbitrary number of control executions can degrade the performance of the safety-critical systems. To ensure performance, *AAwSOS* cannot skip consecutive control tasks beyond a certain limit. In those situations, *AAwSOS* **(c)** will change the ordering of the same priority tasks in e such that none misses the respective deadline to fool the attacker's schedule analysis. We show that schedule obfuscation techniques (a), (b), and (c) reduce the ASP much less than the blind schedule randomization techniques. The primary contributions of this work are summarized below:

(1) We propose an attack-aware schedule obfuscation strategy (*AAwSOS*) *Hide-n-Seek* to *hide* the timing information of safety-critical control tasks transmitted through the CAN bus from an attacker attempting to launch a bus-off attack. The proposed strategy also detects possible bus-off attempts while obfuscating the schedules. It also ensures that a desired performance guarantee is always ensured and none of the tasks misses their deadlines.

(2) We mathematically establish that *AAwSOS* will reduce the attack success probability (ASP) by a significant amount when compared to static schedules or attack-unaware schedule randomization techniques.

(3) We evaluate the efficacy of the proposed *Hide-n-Seek* strategy to combat and detect CAN bus-off attacks on a hardware-in-loop (HIL) setup using benchmark CAN traffic data.

II. BACKGROUND

A. Intra-vehicular Network

Fig. 1: Intra-vehicular Network

We consider a typical automotive closed-loop setup with a CAN network connecting the plants and different ECUs. Such a network is demonstrated in Fig. 1. We consider there are multiple ECUs connected to the CAN bus. Each ECU has multiple cores and a number of tasks (including safety-critical ones) are running on them following certain static schedules. We assume one of the cores of ECUs is always available to run the proposed *Hide-n-Seek* algorithm. The execution of this algorithm is triggered by an interrupt that is raised whenever a message is transmitted in the CAN bus enabling it to observe the CAN traffic. This is considered practical as most of the modern ECUs (eg. Infineon Tricore-397, Raptor GCM-196 etc.) have multi-core implementations. One or more tasks running on ECUs transmit and receive messages to and from the CAN bus. On arrival, such a task instance is executed in the ECU following a static schedule (since the set of tasks on an ECU is fixed), and the data produced by the task is attempted for transmission via CAN as a CAN packet/frame. So the relative order of the ECU tasks and the messages transmitted by them in CAN remains intact.

CAN packets have no source and destination addresses and are communicated in *broadcast* mode. Each data packet has a unique ID that signifies its transmission priority in the bus. Contention to the CAN bus by multiple messages at the same time is resolved through the *CAN arbitration* process. Technically CAN works as a *wired-AND gate*. An ECU sending a message with a higher ID value will see a *dominant bit i.e.* 0 in the bus if another ECU sends a message with a lower ID value at the same time, and thereby, retrieves itself from transmitting. Thus, *a message with a lower ID value is considered to be of higher priority over a message with a higher ID value*. In this work, we focus on the *safety-critical control tasks* that are executed in the ECUs and the control messages transmitted by them via the CAN bus. To understand their importance,

we provide a background on typical automotive control tasks, their modeling, and stability/performance requirements.
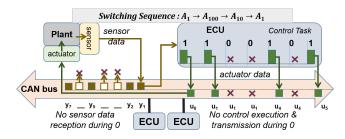

Fig. 2: An embedded CPS following a CSS **1101001**

B. Closed-loop Safety Critical CPS Model

Closed-loop safety-critical automotive CPSs are implemented across multiple ECUs in modern vehicles. Such CPS closed-loops are in charge of different safety-critical real-time operations e.g. anti-lock braking system (ABS), traction control, electronic stability control (ESC), adaptive cruise control (ACC), etc. The controller and the actuator units are implemented in different ECUs that are connected via the CAN bus. The state-space equations for a discrete-time linear-time-invariant (LTI) closed-loop system in Eq. 1 capture the progression of this plant-controller closed-loop. The controller units rely on the sensed measurement data y from the plant, which is communicated through the CAN. Using this an estimator task first estimates the states of the plants. We denote x as the plant's state vector, and $\hat{x}$ as the estimated state vector by the estimator. Based on the estimated plant state, a control task calculates the required control input u in order to transmit it to the actuator for actuation and control of the plant.

$$x[k+1] = Ax[k] + Bu[k], \quad y[k+1] = Cx[k+1]$$
$$\hat{x}[k+1] = (A - LC)\hat{x}[k] + Bu[k] + Ly[k]$$
$$u[k+1] = K\hat{x}[k+1] \tag{1}$$

Here, A, B, C are characteristic matrices. K and L are the controller gain and gain of the Kalman estimator. Indexing the system variables with k simply denotes the sampling iteration of the discrete LTI system. Considering the augmented state vector of the plant-controller closed-loop as $X = [x^T, \hat{x}^T]^T$, and by replacing u with $K\hat{x}$, y with Cx the overall closed-loop system evolves as follows:

$$X[k+1] = A_1 X[k], \ A_1 = \begin{bmatrix} A & BK \\ LC & A - LC + BK \end{bmatrix} \tag{2}$$

We now, introduce the idea of *skipping a control execution* and observe how the aforementioned closed-loop evolves in the presence of such a skip. Skipping a control execution at $(k+1)$-th time instance means the progression of the estimator-controller unit is stalled at that iteration, i.e. $\hat{x}[k+1] = \hat{x}[k]$. Therefore, no sensor data with the output measurement of k-th sampling iteration is transmitted through CAN, and consequently, no new control messages are computed and transmitted at $(k+1)$-th sampling iteration as well. Therefore, at $(k+1)$-th iteration the plant evolves using the last-computed control input i.e. $u[k+1] = u[k]$. In presence of a control

execution skip at $(k+1)$-th iteration, the closed loop system in Eq. 1 evolves as follows.

$$x[k+1] = Ax[k] + Bu[k], \tag{3}$$
$$\hat{x}[k+1] = I\hat{x}[k] + Ou[k] + Oy[k], \quad u[k+1] = K\hat{x}[k+1]$$

I and O are identity matrices and zero matrices with the same dimensions as A, B respectively. The augmented closed-loop system under control execution skip progresses like:

$$X[k+1] = A_0 \ X[k], \ A_0 = \begin{bmatrix} A & BK \\ O & I \end{bmatrix} \tag{4}$$

Due to the occasional control skips, we thus have a switched LTI system where we switch between the different combinations of $\{A_1, A_0\}$. For example, if we execute the control task at k-th iteration and skip the control executions at $(k+1)$-th and $(k+2)$-th iterations, the augmented closed-loop system state at $(k+3)$-th iteration becomes $X[k+3] = A_0 A_0 A_1 X[k]$. This can be represented as a subsystem with state transition matrix $(A_0)^2 A_1$. Switching between such subsystems gives a control execution skip sequence, formally defined below.

*Definition 1 (***Control Skipping Sequence***):* An $l(\in \mathbb{N})$ length control skipping sequence (CSS) for a given control loop is a sequence $\rho \in \{0,1\}^l$ such that 1 denotes a control execution when the augmented state of the plant and controller X progresses following Eq. 2, and 0 denotes a skipped execution when X progresses following Eq. 4. □

A control skipping sub-sequence of an l-length control skipping sequence ρ is an i-length sub-string of ρ having the form 10^{i-1} such that $\sum_q i = l$, where q denotes the number of such sub-sequences in a sequence. A sub-sequence signifies that the plant is actuated with the last control update in consecutive $(i-1)$ iterations. As an example, the control execution sequence $\rho = 1100101$ is generated by 4 times switching between 3 *control skipping sub-sequences* i.e. 1 to 100, 100 to 10 and back to 1 again, i.e. $X[7] = A_1 A_{10} A_{100} A_1 X[0] = (A_0)^0 A_1 (A_0)^1 A_1 (A_0)^2 A_1 (A_0)^0 A_1 X[0]$. Fig. 2 demonstrates the transmission of sensor data and control signals through CAN for an automotive CPS during this control skip sequence.

● ***Performance-aware Switching between Control Skipping Sub-sequences:*** In order to measure the effect of skipping control executions on the performance of the closed-loop system, we represent the augmented closed-loop system as a switched linear system like below.

$$X[k+1] = A_{\sigma[k]}(X[k]) \tag{5}$$

Here, the switching signal $\sigma[k] = i$ signifies that at k-th sampling instance, the system follows the dynamics of the control skipping sub-sequence 10^{i-1}, which we denote as the i-th subsequence. While following an i-th subsequence, a closed-loop skips control execution in consecutive $(i-1)$ iterations and reuses the last computed control update in all these iterations. We use Global Uniform Exponential Stability (GUES) as the performance-driven stability metric to ensure that a CSS maintains a given settling time requirement. The equilibrium at $X = 0$ of a discrete LTI system attains GUES under a certain switching signal $\sigma[k]$ if, for initial conditions

$X[k_0]$, there exist constants $M > 0$, $0 < \gamma < 1$ such that $||X[k]|| \leq M\gamma^{(k-k_0)} ||X[k_0]||$, $\forall k \geq k_0$ where $||.||$ is the vector norm and γ is the GUES decay rate. To ensure that the switched closed-loop system attains certain GUES decay requirement, we *arbitrarily* switch between only those sub-sequences for which, there exists a *Common Lyapunov Function* (CLF). For *arbitrarily switching between such control skipping sub-sequences* the following claim must hold.

Claim 1: A switched system like Eq. 5 can *arbitrarily* switch between a j-th subsequence and an i-th subsequence if and only if a common Lyapunov function (CLF) V exists for all such j-th and i-th subsequences. The CLF should satisfy the following CLF criteria in order to maintain a desired GUES decay margin of γ while switching among them arbitrarily.

$$\kappa_1(||X(k)||) \leq V_i(x(k)) \leq \kappa_2(||X(k)||) \tag{6}$$
$$\Delta V_i(X(k)) \leq \alpha_i \ V_i(x(k)) \text{ and for some } \alpha_i \neq 0 \tag{7}$$
$$V_i(x(k)) \leq \mu_i \ V_j(x(k^-)) \ \forall j, \text{s.t. } i \neq j, \ \mu_i = 1 \tag{8}$$
$$\kappa_1, \kappa_2 \in \text{ class } \mathcal{K}_\infty, \ k^- < k, \sigma(k) = i, \ \sigma(k^-) = j \qquad \square$$

The Lyapunov functions for i-th and j-th subsequences(V_i, V_j respectively), are derived by solving the LMIs for a given desired GUES decay rate γ as mentioned in the above claim. Among such sub-sequences with CLF, the one with maximum number of consecutive control execution skips gives us the upper bound or *skip limit*. For example, consider for a closed-loop, a stable CSS given by $(1 + 10 + 100)^+$ (in regular language notation). The desired GUES decay is generated by arbitrarily switching between the subsequences, 1, 10, and 100. In this example, 2 is the *skip limit* for the closed-loop as maximum two consecutive skips are allowed in the participating sub-sequences. While obfuscating CAN schedules by skipping control tasks, the skip limit helps in preserving control performance of the related closed-loop.

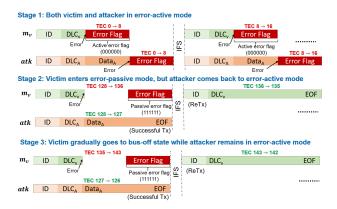### C. Bus-off Attack on CAN: A Schedule-based Attack (SBA)

Fig. 3: Three stages of bus-off attack [4] (m_v: victim message, atk: attack message, EOF: end of frame, IFS: inter-frame space)

In this work, we consider an attacker that attempts to launch a bus-off attack [4] on a victim message ID in the CAN bus. We consider the control messages, originating from the control tasks as the victims in our attack model. By targeting

a control message, the attacker intends *to launch a bus-off attack on the victim ECU* i.e. the source of this control message ID (where the corresponding control task runs). The bus-off attack exploits the error-handling strategy of CAN protocol which states that an ECU (i) remains in *error active* mode if the transmission error count (TEC) or reception error count (REC) is less than 128, (ii) enters an *error passive* mode when TEC or REC exceeds 127, but remains below 256, and (iii) goes to the *bus-off state* i.e. gets disconnected from the bus temporarily when TEC or REC exceeds 255 to notify some significant disruption in transmission. Before launching the bus-off attack, the attacker, which is fundamentally a piece of code implemented on a compromised ECU, targets a safety-critical periodic control message m_v and monitors the CAN bus traffic to compute the timing parameters (periodicity, arrival offsets, etc.) of m_v. It fabricates attack messages atk with the same ID as m_v and a value in the data field less than the one in m_v. Before launching the bus-off attack, it is assumed that both the victim node i.e. the ECU transmitting m_v, and the attacker node, i.e. the compromised/attacker ECU are in error active mode.

The 3 stages of the bus-off attack are demonstrated in Fig. 3. At the first stage, the attacker injects atk in the CAN bus *synchronously* with m_v i.e. exactly at the same time. Since both atk and m_v have the same ID, both win the arbitration process at the ID field. However, during the transmission of the data field, the transmitting ECU of m_v senses the error first, and broadcasts an active error flag (containing 000000). As CAN works as a wired-AND gate, attacker ECU senses transmission error as well. This leads to an increase in TEC by 8 for both the victim and attacker ECUs. Due to the failed transmissions, both victim and attacker ECUs attempt retransmissions of their respective messages at the same time and the same transmission error occurs. After 16 consecutive retransmission attempts, the TEC of both attacker and victim node reach 128 i.e. both the nodes enter error-passive mode. Synchronization of m_v and atk in error-passive mode causes the victim ECU to send a passive error flag (containing 111111) with a further increase in TEC by 8. As this generates no error in the attacker node, the attack message is transmitted successfully by reducing its TEC by 1. As an effect, the victim node's subsequent retransmission attempt becomes successful without any obstacle from the attacker and reduces its TEC by 1. In the second stage of the attack, while the attacker is in error passive mode, it again synchronizes with the victim during the next transmission cycle. The attacker's TEC keeps on reducing by 1, and the victim's TEC increases by $8 - 1 = 7$. At this third stage of the attack, after subsequent synchronous transmissions in error passive mode, the attacker comes back to the error active mode but the victim eventually goes into the bus-off state. In the next sections, we present our novel schedule obfuscation strategy *AAwSOS* utilizing control execution skips. Our *AAwSOS* can *hide* the CAN bus schedule of a victim controller message and prevent the bus-off attack attempt from becoming successful while *seeking* to detect the attacker.

III. Proposed Methodology

Before we explain our attack-aware schedule obfuscation policies, we need to understand how a bus-off attacker looks for the *oppertunest* moments to launch a bus-off attack. This serves as a motivation behind the innovation of various schedule obfuscation policies in an attack-aware way.

A. Schedule Analysis To Launch Bus-off Attack

The success of the bus-off attack heavily relies on how precisely the attack message atk synchronizes with the targeted instance of victim message m_v (see Fig. 3). The attacker analyzes the CAN schedule to find out targeting which instance of m_v results in the highest probability of synchronization, and thereby, bus-off attack [1]. As shown in Fig. 4, there can be 3 types of message-sequences in CAN bus preceding each instance of m_v: i) one or more higher-priority messages, i.e. whose IDs are lower than m_v's (e.g., 3-rd and 4-th instances of m_v in the 1-st hyper-period of CAN bus in Fig. 4); ii) one or more lower-priority messages, i.e. whose IDs are higher than m_v's (e.g., the 2-nd instance of m_v in the 1-st hyper-period of CAN bus in Fig. 4), and iii) idle time slot (e.g., the 1-st instance of m_v in the 1-st hyper-period of CAN bus in Fig. 4). In the case of (iii), the attacker has to inject atk at exactly the same time as m_v to initiate the bus-off attack. Perfect synchronization of atk with m_v will be extremely difficult due to bus jitters. Therefore, the attack success probability (ASP) is nearly negligible in this scenario. In the case of (ii), the attack initiation window is, at most, the transmission time of the immediately preceding lower-priority message. Note that the attacker must precisely inject atk after the lower-priority message wins the bus arbitration, i.e. after the transmission of the ID field of the lower-priority message. Otherwise, atk will win the arbitration as it has higher priority. This makes the atk face the CAN arbitration before m_v and hence loses the intended synchronization with m_v. In the case of (i), the attacker can initiate the bus-off attack anytime during the transmission of one or more consecutive higher-priority messages preceding m_v. This will force atk to synchronize with m_v after the transmission of the last higher-priority message preceding m_v. Therefore, the ASP, in this case, is much higher than (ii) and (iii).

In this schedule-based attack (SBA) setup, we consider the following two assumptions regarding the attack model. *First*, the attacker targets a messages m_v corresponding to a safety-critical control task T_v. *Second*, it will always target those instances of an m_v that immediately follow one or more higher-priority messages in the bus (i.e. type (i) message sequence mentioned in the previous paragraph). Therefore, we consider this sequence of high-priority tasks coming just before the instances of a victim message m_v in the CAN bus as *attack window*. The ASP corresponding to each instance of m_v changes proportionally with the attack window length. The attacker observes the CAN schedule for a number of hyper-periods, say $recon$, to identify which instance of the victim message can be targeted for high ASP. We call this duration as *reconnaissance period*. The attacker computes the

average attack window length for each instance of m_v over $recon$ CAN hyper-periods. The ASP will be highest if the attacker targets the particular instance of m_v where the attack window is the longest. For example, in Fig. 4, the average numbers of higher-priority messages coming just before the 4 instances of m_v over $recon = 2$ hyper-periods in CAN bus are $\lceil (0+1)/2 \rceil = 1$, $\lceil (0+2)/2 \rceil = 1$, $\lceil (3+3)/2 \rceil = 3$, and $\lceil (5+5)/2 \rceil = 5$ respectively. Therefore, targeting the 4-th instance of m_v in the next reconnaissance period will lead to maximum ASP for an attacker.

B. Attack-aware Schedule Obfuscation Strategies

The primary motivation of this work is to devise a lightweight strategy to break the attacker's schedule analysis for finding the target instance of victim message m_v that exhibits high ASP. In this regard, we propose 3 attack-aware schedule obfuscation strategies (*AAwSOS*). Based on the analysis of the CAN traffic during the current reconnaissance period, the *Hide-n-Seek* framework implemented in each ECU will instrument a task execution schedule for the corresponding ECU during the subsequent reconnaissance period. The objective of the *Hide-n-Seek* algorithm is to obfuscate the future control task schedules in an ECU to introduce a maximum reduction in ASP of the resulting global CAN schedule. In Fig. 4, we provide a snapshot of the CAN bus in the current reconnaissance period whose length is taken as 2 consecutive hyper-periods. By analyzing this CAN message schedule in the current reconnaissance period, *Hide-n-Seek* obfuscates the existing control task execution schedule that is to be followed during the next reconnaissance period as demonstrated in Fig. 5. This obfuscation decision is not a static one, as the CAN traffic changes due to (i) the local obfuscation decisions taken by the other ECUs connected to the bus, (ii) timing interference caused by the arrivals of 'non-actuation' and aperiodic messages, and (iii) obviously, due to the undesired delays and jitters. This demands a cyclical execution of our *Hide-n-Seek* algorithm for every reconnaissance period. The idea is to observe the current CAN traffic and instrument the next ECU schedule for every sliding window of the length same as the reconnaissance period. While doing so, the *Hide-n-Seek* algorithm observes the bus and flags the message instances corresponding to a scheduled skip in task execution and thereby detects a bus-off attack attempt.

● **Obfuscation Strategy 1** (Obf_1) - *Skipping the vicim task instance*: As discussed in Sec. II-A, the relative orders of the tasks executed in the ECU's schedule and their corresponding messages being transmitted through CAN, are the same. The schedule analyzer implemented in each ECU analyzes the CAN bus for a reconnaissance period of $recon$ CAN hyper-periods and sorts the instances of target message m_v of each safety-critical task T_v running in that ECU in descending order of corresponding ASP. Following this order, the Obf_1 will skip the corresponding instances of T_v one by one in the ECU's schedule as long as the CLF property holds (Claim. 1). For example, by observing the CAN schedule for $recon = 2$ hyper-periods in Fig. 4, we can see the 4-th instance of m_v in
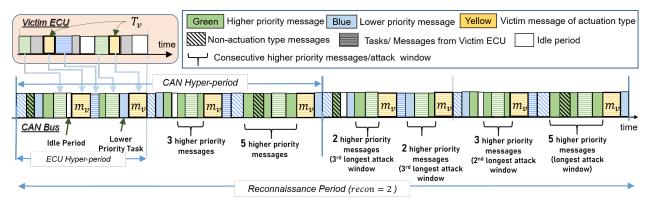
Fig. 4: CAN schedule analysis for launching bus-off [the green/blue 'unpatterned' messages signify control tasks in other than victim ECUs]
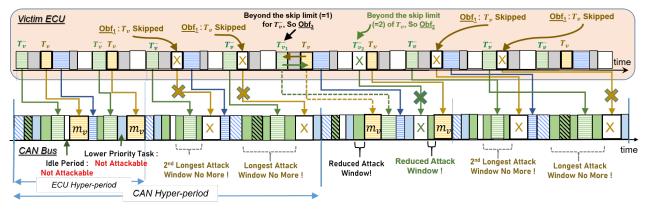

Fig. 5: CAN schedule analysis for launching bus-off Under *Hide-n-Seek* [legends are same as Fig. 4]

a CAN hyper-period exhibits the highest ASP followed by 3-rd, 2-nd, and 1-st instances of m_v. Accordingly, Obf_1 decides to skip the 4-th and 3-rd instances of T_v in the subsequent schedules of the victim ECU as shown in Fig. 5. As an effect, at each hyper-period of the next reconnaissance period, the 4-th and 3-rd instances of m_v will not be present in the bus.

• **Obfuscation Strategy 2** (Obf_2) - *Skipping instances of the tasks with a priority higher than victim task*: Obf_1 cannot skip consecutive task instances beyond the *skip limit* set by the CLF property of the corresponding controller (see Sec. II-B). For example, the skip limit of the control task T_v in Fig. 5 is 2 according to its CLF property. Therefore, Obf_1 cannot choose to skip the 2-nd instance of T_v. Since these strategies are to be chosen locally in an ECU, it can only decide which among the tasks executing on the current ECU can be skipped to reduce ASP. If instances of T_v cannot be skipped to respect the skip limit, as part of Obfuscation Strategy 2, the *Hide-n-Seek* module skips a control task that (i) is of higher priority than T_v, (ii) transmits a control message with ID less than m_v, and (iii) contributes to the attack window before the current instance of m_v. To elaborate, if T_v^- is the higher priority task that executes before T_v following the static schedule in the victim ECU, its corresponding message m_v^- comes before m_v in the CAN bus. If m_v^- falls inside the attack window of m_v, skipping T_v^- will reduce m_v's attack window length. This reduces the ASP of a targeted message instance. In the best case, the ASP may even reduce to 0. For example, notice the

attack window of the 2-nd instance of m_v in the 2-nd CAN hyper-period in Fig. 4. One of the higher-priority messages in this attack window corresponds to a task in the victim ECU having higher priority than T_v. Skipping this higher priority task in the ECU's schedule will reduce the attack window size to 1 in the next reconnaissance period as shown in Fig. 5, and thereby, reduce the corresponding ASP.

• **Obfuscation Strategy 3** (Obf_3) - *Executing the victim task first among the same priority tasks*: As discussed earlier, any control task cannot be skipped beyond the skip limit set by its CLF property. For example, Obf_2 policy has skipped executing one instance of T_v^-, say $T_{v_2}^-$, in victim ECU whose corresponding message instance m_v^- belongs to the attack window of m_v's 2-nd instance. Considering the skip limit for T_v^- is 1, Obf_2 cannot skip the particular instance of T_v^-, say $T_{v_1}^-$, whose corresponding message instance m_v^- falls into the attack window of m_v's 1-st instance (see the 2-nd CAN hyper-period in Fig. 5). As can be seen in Fig. 5, the two instances $T_{v_1}^-$ and $T_{v_2}^-$ of T_v^- come sequentially in the ECU's schedule. Either of $T_{v_1}^-$ and $T_{v_2}^-$ can be skipped, but not both, according to T_v^-'s CLF property. In such a scenario, the *Hide-n-Seek* framework opts for another obfuscation strategy Obf_3. Obf_3 figures out the set of tasks $T_v^<$ (includes T_v) in the ECU i) having the same priority as T_v, and also, ii) belonging to the attack window of m_v in the CAN bus. At certain time instances when the scheduler in ECU has to pick one among the tasks in $T_v^<$, Obf_3 dictates it to select the victim task T_v

first during the subsequent reconnaissance period. This ensures that the set of higher priority messages $m_{v<}$ corresponding to the tasks in $T_v^<$ will come after m_v in the next reconnaissance period. If the cardinality of the set $T_v^<$ (and $m_v^<$) is $|T_v^<|$ ($|m_v^<|$), the application of obfuscation strategy 3 will reduce the attack window of m_v by $|m_v^<|$, and thereby, reduce the corresponding ASP. In Fig. 5, consider that the yellow and green patterned and the grey tasks have the same priority in the victim ECU. However, only the message corresponding to the green patterned task belongs to the attack window of the 1-st instance of m_v (in the 2-nd CAN hyper-period). We can see that executing T_v before the green patterned task in the ECU's schedule reduces the attack window of m_v's 1-st instance from 2 to 1.

AAwSOS - **Defends as well as Detects SBA attempts**: Analyzing the CAN traffic during the current reconnaissance period, the attacker finds out the target instances of the victim message with higher ASP i.e. non-zero attack window length. The obfuscation strategies discussed above ensure the reduction of ASP of each such probable target message instance by *hiding* their timing information in the next reconnaissance period. The *AAwSOS* strategies either skip or relocate the target message instances. An attacker may attempt to launch the SBA in the next reconnaissance period by injecting a message with the same ID as the victim message during the attack window. The *AAwSOS* strategies invalidate the attacker's calculation of the precise arrival time of the target message instances, and thereby, quash the SBA attempt. Consequently, *AAwSOS* strategies help in *seeking* the attack attempt. For example, analyzing the current reconnaissance period of $recon$=2 CAN hyper-periods in Fig. 4, the attacker may target the 4-th instance of m_v to launch SBA. It fabricates atk message with the same ID as m_v and injects it into the CAN bus during the attack window of m_v's 4-th instance in the next reconnaissance period. However, since Obf_1 has skipped the corresponding execution of T_v in the victim ECU, ideally the 4-th instance of m_v should not be present in the CAN bus. Observing a message with the same ID as m_v in the CAN bus during the same time slot, the *Hide-n-Seek* module will perceive it as coming from a compromised ECU and raise an alarm to indicate a bus-off attempt.

C. Attack Success Probability under Hide-n-Seek

To analyze how effectively the proposed *AAwSOS* strategy defends against SBAs, we evaluate the attack success probability (ASP) on a CAN hyper-period with and without this strategy. As explained in Sec. III-A, launching an SBA requires a detailed timing analysis of CAN traffic for a certain number of hyper-periods which we call *reconnaissance period*. Let us consider that the reconnaissance period for an attacker is $recon$ number of CAN hyper-periods (i.e. $\{H_1, H_2 \cdots H_{recon}\}$). Based on the CAN transmission schedule, we divide each CAN hyper-period into multiple slots. Consider there are n_i transmissions with ID i for a task T_i where $i \in IDset$. Here, $IDset$ denotes the set of message IDs corresponding to control tasks transmitted over every CAN hyper-period

of length H. This means the message with ID i uses n_i number of slots for transmissions in H duration. If the j-th slot in H is used to transmit the message with ID i, then the length of the j-th slot equals the time required to transmit the message ID i. The total number of slots for transmission of all message IDs in a single CAN hyper-period is $J = \sum_{\forall i \in IDset} n_i$. Now, to quantify an SBA effort, we utilize this concept of slots. Let us consider T_v as the victim task that transmits the message m_v with ID v. As explained in Sec. III-A, there are two desired events that the attacker observes during the reconnaissance period for launching a successful SBA. *First*, the attacker needs to check whether the victim message m_v arrives at j-th slot of every hyper-period during reconnaissance (i.e. during $\{H_1, H_2, \cdots, H_{recon}\}$). *Next*, the attacker observes whether there is *non-zero* number of higher priority messages before the victim message m_v at every hyper-period during reconnaissance. This helps the attacker decide whether or not to launch SBA on m_v at j-th slot. We define a function $ct : IDset \times [1, n] \mapsto 0, 1$ as follows.$ct(v, j) = 1$if T_v transmits (m_v) at j-th slot and $ct = 0$if T_v does not transmit at j-th slot. Let us consider that the number of higher priority messages preceding the target/victim ID v at j-th slot is $n_{v,j}$. To quantify the chance of launching a successful SBA at j-th slot, we can measure the ratio between the number of slots/transmissions with higher priority and the total number of transmissions between the current and last instance of m_v. If we denote the count of the total number of **t**ransmissions **b**etween the current and last **i**nstances of m_v at j-th slot as $TBI_{v,j}$ the *attack success* at j-th slot can be measured as, $P(AS_j) = \sum_{k=1}^{recon} \frac{ct(v,kJ+j)}{recon} \cdot \frac{n_{v,kJ+j}}{TBI_{v,kJ+j}}$. This is nothing but the average of the ratios of higher priority tasks before the victim ID to TBI ($\frac{n_{v,j}}{TBI_{v,j}}$) at every j-th slot where it appears in every hyper-period $\in \{H_1, H_2, \cdots, H_{recon}\}$.

A successful SBA launch event on a victim ID m_v occurs when SBA can be launched at any one of the J transmission slots during the reconnaissance period. Since being successful at different slots are mutually exclusive events, the overall SBA success probability for the attacker becomes $P(AS) = P(\bigcup_{\forall j \in [1, recon.J]} AS_j) = \sum_{\forall j \in [1, recon.J]} P(AS_j)$. To establish the effectiveness of our obfuscation policy, we claim the following.

Claim 2: Application of the *Attack-aware schedule obfuscation strategies* (Sec. III-B) reduces the *ASP* of SBA. $\square$

Proof: As discussed in Sec. III-B, *Hide-n-Seek* has *three* obfuscation strategies $\{Obf_1, Obf_2, Obf_3\}$, either of which we apply at every victim task instance in the victim ECU i.e. at every j-th slot of a CAN hyper-period, whenever the victim task m_v appears. Let us denote O_i as the event of applying obfuscation strategy Obf_i where $i \in \{1, 2, 3\}$. It is intuitive that O_i's are mutually exclusive. Therefore, the probability of applying any obfuscation strategy can be calculated $P(Obf) = P(O_1 \cup O_2 \cup O_3) = \sum_{i \in 1,2,3} P(O_i)$. The ASP under *Hide-n-Seek* can therefore be evaluated using the

following conditional probability.

$$P(AS|Obf) = P\left(\bigcup_{\forall j \in [1,reconJ]} (AS_j|(O_1 \cup O_2 \cup O_3))\right)$$

$$= \sum_{\substack{\forall j \in [1,reconJ] \\ s.t. \ ct(v,j)=1}} \frac{(P(AS_j \cap (O_1 \cup O_2 \cup O_3))}{P(O_1 \cup O_2 \cup O_3)}$$

$$= \sum_{\substack{\forall j \in [1,reconJ] \\ s.t. \ ct(v,j)=1}} \frac{\sum\limits_{\forall q \in \{1,2,3\}} P(AS_j|O_q)P(O_q)}{\sum\limits_{\forall q \in \{1,2,3\}} P(O_q)}$$

If we apply either of the obfuscation policies at every instance of the victim task execution in the victim ECU which transmits the victim message m_v in CAN, then the denominator of the above conditional is evaluated to be 1 (i.e. O_1, O_2 and O_3 are exhaustive). Hence,

$$P(AS|Obf) \approx \sum_{\substack{\forall j \in [1,reconJ] \\ s.t. \ ct(v,j)=1}} \sum_{q \in \{1,2,3\}} P(AS_j|O_q)P(O_q)$$

$$= \sum_{\substack{\forall j \in [1,reconJ] \\ s.t. \ ct(v,j)=1}} P(AS_j|Obf)$$

$$\Rightarrow P(AS_j|Obf) \approx \sum_{q \in \{1,2,3\}} P(AS_j|O_q)P(O_q)$$

Now, let us observe how the probability of the *attack event* under the application of the obfuscation rules, changes.

- The Obf_1 strategy skips the execution of the T_v instance which transmits m_v at j_v-th slot. Therefore, we can say on application of Obf_1, $ct(v,j_v) = 0 \Rightarrow P(AS_{j_v}|O_1) = 0 < P(AS_{j_v})$.

- Similarly, as per the Obf_2 strategy, the execution of the higher-priority task instance T_v^- is skipped as it is executed before the current instance of T_v and contributes to the attack window before m_v at j_v-th slot. In the best case, the attack window length can become 0. Therefore, the application of Obf_2 reduces the attack window length from n_{v,j_v} to n_{v,j_v}' such that $n_{v,j_v}' < n_{v,j_v}$. This implies $P(AS_{j_v}|O_2) < P(AS_{j_v})$.

- Finally, Obf_3 strategy shifts the execution of the job-instances of T_v which transmits m_v at j_v-th slot, before all the tasks of its same priority. We do so, only if the same-priority tasks in ECU-level transmit higher priority messages in CAN that contribute to the attack window before m_v at j_v-th slot. Therefore, we can say the application of Obf_3 reduces the number of preceding higher-priority messages. Say, the set of tasks that are of the same priority as T_v and have contributed to the attack window before m_v in CAN's j_v-th slot, is denoted as $T_v^<$. The Obf_3 will reduce the attack window length before m_v to $n_{v,j_v}' = n_{v,j_v} - |T_v^<| < n_{v,j_v}$. This implies $P(AS_{j_v}|O_3) < P(AS_{j_v})$. As the *attack event conditional under obfuscation policies* is less than the actual attack event probabilities without obfuscations, the ASP at any j-th slot is definitely less when the obfuscation strategies are applied, i.e. $P(AS_j|Obf) < P(AS_j)$. Since the overall ASP is the sum of ASPs at every j-th slot where the victim task arrives within the reconnaissance period, the overall ASP under our

AAwSOS is also reduced. This proves the claim. □

- **Comparison with *typical schedule randomization* Strategies:** Notice that the attack event conditional under *Hide-n-Seek* policy is nothing but a weighted average of attack success conditionals under different obfuscation policies: $P(AS_j|Obf) = \sum\limits_{q \in \{1,2,3\}} P(AS_j|O_q)P(O_q)$. As part of *Hide-n-Seek*, we apply Obf_1 and Obf_2 as our primary preferences. As discussed earlier, $P(AS_j|O_1) = 0$ and $P(AS_j|O_2) < P(AS_j)$. When we do not have the luxury to skip any more control executions in consecutive iterations to respect the *skip-upper-bounds* of the closed-loops, we apply Obf_3 as our third choice. This means, $P(O_3)$ depends on the upper bounds of the consecutive execution skips. Note that, $P(AS_j)$ due to the application of Obf_3 policy is $\frac{n_{v,j_v} - |T_v^<|}{TBI_j}$. This leads to a reduction in the ASP by $\frac{|T_v^<|}{TBI_j}$ from the situation before the schedule transformation with Obf_3.

Attack-unaware schedule randomizations have been proposed in earlier works in order to conceal timing side channels exposed by processor-level task schedules [11]. Note that w.r.t. a bus-off attacker, the most resilient schedule that may come out of such non-informed schedule randomization is basically the outcome of the application of Obf_3. In this case, task skipping does not happen, so Obf_1 and Obf_2 do not exist. Hence, under randomization techniques like [11], the ASP lower bound is $\frac{n_{v,j} - |T_v^<|}{TBI_j}$. Whereas, in the case of the *AAwSOS*, the ASP is significantly less due to additional application of the other two policies (i.e. Obf_1 and Obf_2) that have the ability to reduce $P(AS_j|Obf)$ to $zero$ in suitable cases. This evidently concludes the fact that our *AAwSOS* policy is more apt in terms of defending against SBAs compared to any other schedule randomization technique.

D. Algorithmic Framework for Hide-n-Seek

We propose an algorithmic framework of the proposed *Hide-n-Seek* strategy that should be locally implemented on every ECU connected to the CAN bus. We assume the execution of the task set TS follows a static schedule in an ECU. As a prerequisite to running the algorithm, we consider that a thorough analysis of CAN traffic for $recon$ number of CAN hyper-periods is done following the state-of-the-art [1] methodologies. Following is the information available from this analysis: (i) the length of CAN hyper-period H and the minimum number of CAN hyper-periods required for reconnaissance $recon$; (ii) the set of control message IDs transmitted via the CAN bus $IDset$; (iii) The ID-wise list of attack windows (containing the consecutive higher priority transmission slots) appearing in front of any instance of a message ID generated from the current ECU ($atkWinlist$). With these as inputs, in Algo. 1, we make two procedure calls (i) HIDE() that decide the obfuscation policies for each of the safety-critical control task instances running in the current ECU (line 3); (ii) SEEK() that seeks for the presence of a bus-off attacker in CAN bus (line 4).

In each call for every reconnaissance period, the HIDE() procedure is run for rpt number of ECU hyper-periods

```
Require: H, recon, IDset, atkW inlist, TS, h, skiplist
Ensure: ObfScheds, alarm
 1: function HIDE-N-SEEK(H, recon, IDset, atkW inlist, TS, h, skiplist)
 2:    for every reconnaissance period do
 3:       ObfScheds ←HIDE(H, recon, IDset, atkW inlist, TS, h,
    skiplist)
 4:       alarm ← SEEK(H, recon, ObfScheds, skiplist)
 5:       if alarm then take safety/security measure
 6:    return ObfScheds, alarm
 7: function HIDE(H, recon, IDset, atkW inlist, TS, h, skiplist)
 8:    rpt ← ⌈ recon×H / h ⌉
 9:    Scheds, ObfScheds, eqPrilists ← [ ], [ ], [ ]
10:    while rpt > 0 do
11:       eqPrilists, Sched ← Scheduler(TS)
12:       Scheds ← Scheds.append(Sched)
13:       eqPrilists ← eqPrilists.append(eqPrilist)
14:       ObfScheds ← ObfScheds.append(Sched)
15:       rpt ← rpt − 1
16:    sort IDset in descending order of attack window length in atkW ins
17:    for each id ∈ IDset do
18:       atkW ins ← atkW inlist[id]
19:       ideqPrilist ← eqPrilists[id]
20:       sort atkW ins in descending order of slot vulnerability value
21:       ECUslots ← slotMapper(atkW ins.slots)
22:       for each si ∈ ECUslots do
23:          ifObf₁ ← CHECKSKIPLIM(id, si, ObfScheds, skiplist)
24:          if Obf₁ then ObfScheds[si] ← 0
25:          else
26:             idPreds, sPreds ← GETHPPREDS(ideqPrilist[id])
27:             sort idPreds, sPreds in descending order of start time
28:             idPred, sPred ← idPreds[0], sPreds[idPreds[0]]
29:             for each idPred ∈ idPreds do
30:                ifObf₂ ← CHECKSKIPLIM(idPred, ObfScheds, skiplist)
31:                sPred ← sPreds[idPred]
32:                if ifObf₂ then ObfScheds[sPred] ← 0
33:                   break
34:             if ¬(ifObf₁ ∧ ifObf₂) then                    ▷ Obf₃
35:                ObfScheds ← SHIFTVIC(ObfScheds, si, sPreds)
36:    return ObfScheds
37: function SEEK(H, recon, ObfScheds, skiplist)
38:    si, alarm ← 0, 0
39:    for si < recon.J do
40:       ECUSlot ← slotMapper(si)
41:       if ObfScheds[si] ≠ ECUSlot then
42:          alarm ← 1
       return alarm
```

Algorithm 1: Control-skip-guided Attack-aware Schedule Obfuscation and Detection Algorithms

(lines 7-36). Here rpt is calculated in order to cover the whole reconnaissance period i.e. $recon \times H$ (see line 8). Since the task set TS is static, by observing the task schedule in the current hyper-periods, we can derive the task execution schedule for the coming rpt number of ECU hyper-periods. We store this extended schedule in $Scheds$ (line 12). While deriving this schedule, we also find out the list of equal priority tasks available with every control task instance at the same time in the ECU processor. This is stored in $eqPrilists$ as shown in line 13. An array $ObfScheds$ is created to store the obfuscated schedules for the next reconnaissance period and initialized with the derived extended schedule (line 14). Next, we sort the $IDset$ in descending order of the attack window lengths of each ID as stored in $atkW inlist$ (see line 16). This enables us to decide the obfuscation policies for all the control message IDs, starting from the most vulnerable ID in the $IDset$ (in line 17-36). Note that, ASP increases with attack window length (Sec. III-A).

We collect the lists of attack windows and the list of equal priority task information in $atkW ins$ and $ideqPrilist$ respectively for each victim ID in $IDset$ (see lines 18, 19). In

line 20, the transmission instances of the current ID, stored in $atkW ins$, are sorted in descending order of the attack window length. Thereafter, we find the ECU task execution slots corresponding to these sorted CAN transmission instances which are then stored in $Ecuslots$ maintaining the order (refer to line 21). A subroutine SLOTMAPPER is used here that observes the relative order of the message transmissions through the CAN bus and finds out the corresponding task execution timeslot in the ECU. Having these $ECUslots$ in the descending order of attack window length, we then decide which policy to apply on which one in lines 22-35.

As discussed in Sec. III-B, we primarily decide to apply Obf_1 at the most attackable task instance. For this, we must check whether skipping at the current task execution slot does not violate the mandated skip limit for the control task. We check this using a subroutine CHECKSKIPLIM() that takes the task ID, the skip limit of the task ID from $skiplist$ (*skip limits* of every control task running in this ECU is stored in $skiplist$), and the obfuscated schedule as inputs and returns a boolean decision to denote whether a skip at the current time slot is permitted for the input task ID. For a certain ECU timeslot of the current task, we store the result in a boolean variable $ifObf_1$ in line 23. When $ifObf_1$ is $true$, we skip the task execution in the current timeslot and store a 0 to denote an idle period at the current timeslot of $ObfScheds$ (line 24). If skipping at the current task instance is not permissible then we go for Obf_2. The $ideqPrilist$ contains the list of tasks that are of equal priority at the currently considered execution slot of a task as per the ECU schedule. A subroutine GETHPPREDS() is designed that takes this list as input and returns a list of tasks (along with their corresponding execution slots) that contributes to the attack window of the message corresponding to the currently considered victim task instance. These IDs and their execution slots are stored in $idPreds$ and $sPreds$ for further analysis (refer to line 26). These lists are sorted in descending order of the start time of the execution slots (in line 27). This helps us start the analysis for the application of Obf_2 with the task that transmits a higher priority message at the nearest slot before the current instance of the victim message. Starting from this, for every task in the $idPreds$ set, we check whether the skip limit of that task supports an execution skip at the current time slot $sPred$ (see line 29). This decision is updated in the decision variable Obf_2 in line 30. On finding such a task, we stop checking further in $idPreds$ and skip the execution of the chosen task at the $sPred$ slot by storing a 0 in the $sPred$ slot of the obfuscated schedule array $ObfScheds$ (see line 32). When both obfuscation policies Obf_1 and Obf_2 are not applicable, we apply Obf_3 on the current task schedule (see line 34). The subroutine SHIFTVIC is designed to update the schedule with a valid schedule following certain scheduling algorithms (eg. EDF, RM) such that the attack window size is reduced. We apply Obf_3 and update the obfuscated schedule array $ObfSched$ (see line 35). After finalizing the strategies for each of the instances of each ID, we finally return the obfuscated schedule as output in line 36.

The SEEK() procedure is designed to raise an alarm in the presence of an attacker in the CAN bus (lines 37-42). We run this method for all ECU task execution slots spanning through the reconnaissance window (see line 39). As shown in line 40, we first check whether a control task is transmitting in the current slot. This finding is then verified against the obfuscated schedule of its corresponding task execution slot (line 41). If an instance of a control message ID appears in the bus while the decided obfuscation strategy for this message instance, mandates skipping of the corresponding task instance, then we set and return the $alarm$ to notify a schedule-based attack attempt (lines 42).

The two procedures HIDE() and SEEK() are deployed in each ECU and run for every reconnaissance period (lines 2-5). The HIDE() procedure uses the data analyzed from the CAN traffic in the current reconnaissance period and decides the obfuscation strategies for the ECU schedules of the next reconnaissance period. It returns the obfuscated schedule $ObfScheds$ (line 3) following which the tasks are scheduled in ECU during the next reconnaissance period. On finding a message corresponding to a skipped task execution, the SEEK() procedure raises an alarm notifying a bus-off attack attempt (line 4). The system is designed to take certain safety and security measures in such a situation in order to refrain from any adverse effects of the attack (line 5). This cyclically goes on for all the reconnaissance periods.

IV. EXPERIMENTAL RESULTS

We demonstrate the efficacy of the proposed $Hide$-n-$Seek$ algorithm on practical CAN traffic under different busloads. For this, we consider a set of control tasks that are implemented across two ECUs. We use Infineon Tricore-397 and Tricore-234 ECUs for this purpose. An automotive plant is emulated in a HIL setup which is in a closed loop with these control tasks. ETAS Labcar real-time simulator is used for this purpose. The plant and the controllers communicate via the CAN bus operating with 250kbps bandwidth. The control tasks and the $Hide$-n-$Seek$ algorithm are implemented and run in different cores of ECUs. A CAN schedule analysis algorithm is run in the same core as the $Hide$-n-$Seek$ algorithm in order to observe all transmissions through the CAN bus. We enlist the detailed specifications of the tasks running in each of the ECUs and their corresponding messages in Tab. I.

TABLE I: Task Setup

ECU	Task Type	Message ID
ECU 1 (TC397)	Control Tasks	0xB0, 0xC0 0xC4
	Other Tasks (* aperiodic tasks)	0xA0, 0xB1, 0xA3, 0xB3, 0x163*
ECU 2 (TC234)	Control Tasks	0x230, 0x323 0x90
	Other Tasks (* aperiodic tasks)	0x98, 0x500, 0xB2, 0xC2, 0xB5, 0xC5, 0x565*, 0x783*

We schedule the periodic control tasks in their corresponding ECUs following an EDF schedule with 60 ms and 30 ms ECU hyper-periods. Whereas the CAN hyper-period is observed to be $\approx$ 200ms. We consider a 3 CAN hyper period long reconnaissance period of length $\approx$ 600ms. The proposed $Hide$-n-$Seek$ algorithm runs for all control tasks in $ECU1$. As part of the schedule analysis, a CAN data logger is implemented to observe all messages as they appear in the CAN bus to estimate the possible attack windows for the control tasks running in this ECU during the reconnaissance period. Based on these estimations, $Hide$-n-$Seek$ decides which obfuscation rules to apply on which instances of the ECU task.

• *Evaluation of the Proposed Obfuscation Policies:* Fig. 6 demonstrates such observed ASPs at every task instance when the tasks are scheduled with a static EDF schedule (the $circled$ $green$ plots) and the resulting ASPs on the deployment of suitable policies at every task instance in the subsequent reconnaissance period (red plots with $cross$ markers). The x-axis in Fig. 6 denotes the relative slot index in every reconnaissance period and their instance-wise ASPs at the y-axis and their corresponding relative slot index in the x-axis. We plot the ASP of the current reconnaissance period and the next reconnaissance period in the same plot for comparative visualization. The reductions in ASPs caused by $AAwSOS$ application can be easily gauged from the plots. The success of a bus-off attack highly relies on the traffic load in the bus, because with more traffic there can be longer attack windows and a better chance of going to bus-off from error passive mode due to passive error regeneration [15]. Therefore we report our results for 3 possible busload conditions i.e. under high (75%), medium (55%), and low (25%) busload. As can be seen in Fig. 6g, Obf_1 is applied at those execution instances of task T_{B0} (the task responsible for the transmission of message ID 0xB0) in the 2nd reconnaissance period where maximum ASPs (≈ 0.12) are observed under a static task schedule during the 1st reconnaissance period (i.e. at the execution slot number $2, 5, 6, 9, 10, 13$ and 14 for T_{B0}). This results in zero ASPs at their corresponding CAN transmission slots and leads to detection of any bus-off attack attempted during those periods (Notice in Fig. 6g alarm flag is raised at 5-th slot). At 11-th instance of T_{B0} as Obf_1 cannot be applied in order to respect the skip limit, Obf_2 is applied by skipping execution of another control task T_{C0} (see at 11-th task instance in Fig. 6h) that contributes to the attack window of T_{B0}.

At certain execution instances, $Hide$-n-$Seek$ applies Obf_3 and reduces the ASPs where the Obf_1 policy can not be applied to avoid violating the skip limit of T_{B0} and Obf_2 policy cannot be applied as none of the preceding control tasks executes and transmits messages during the attack window of T_{B0} (see the 3rd, 7th and 15th instances). As can be seen in Fig. 6i, in case of T_{C4}, $Hide$-n-$Seek$ applies execution skips (i.e. Obf_1) almost at alternate instances and is able to reduce the ASP to zero at every transmission instance of ID 0xC4. In the presence of a medium busload, the reduction in ASP is not this significant but the possibility of detection persists (see 6f). It is observable from the plots that the ASPs reduce as the busload decreases. However, it is also evident, that our obfuscation strategy reduces the ASPs under any busload.

• *Detection using* Hide-n-Seek : To demonstrate the fact that our *AAwSOS* applies the obfuscation rules without hampering the performances of any control loops, we demonstrate the performance of the TTC closed loop under the obfuscations applied on the control task T_{C4} (which transmits with ID
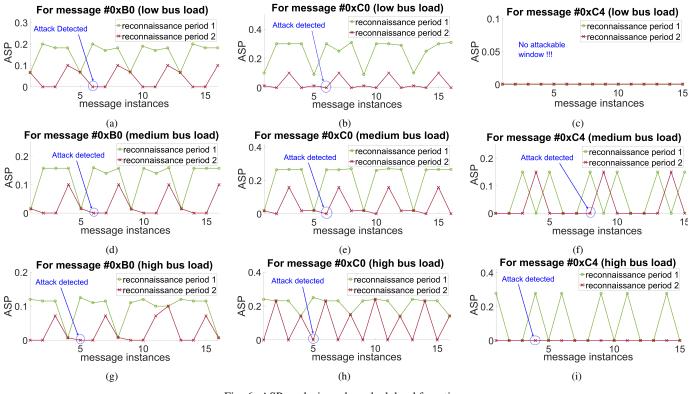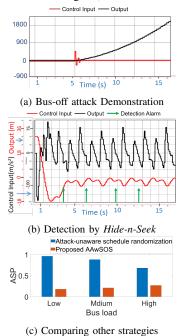
Fig. 6: ASP analysis under schedule obfuscation.

0xC4) during high busload situation (obfuscation decisions are shown in Fig. 6i).

(a) Bus-off attack Demonstration

(b) Detection by *Hide-n-Seek*

(c) Comparing other strategies

Fig. 7: Detection and Comparison.

In the absence of any attack-aware schedule obfuscation, a bus-off attempt on the victim message ID 0xC4 is successful (see Fig. 7a). On the other hand, the presence of *AAwSOS* mandates a task execution schedule with control skips introduced at alternate instances (which respects the skip limit 1 for TTC).

Moreover, the bus-off attack attempts are detected during the applications of Obf_1 (marked with the green arrow in Fig.7b).

● *Comparison with Attack-unaware Schedule Randomization:* To quantify the effectiveness of our *AAwSOS* compared to state-of-the-art schedule randomization techniques, we run a similar set of ASP analysis experiments (for the scenarios shown Fig. 6) applying typical schedule randomization techniques against a bus-off attempt. As discussed in Sec. III-C, such policies are unable to reduce the total ASP ($P(AS)$) as they blindly randomize the

task execution schedules. As can be seen in Fig. 7c, such randomization policies showcase a high ASP in the presence of a low busload (the blue bar plots in Fig. 7c). The ASP under the presence of schedule randomization reduces as the busload increase. When compared to the ASPs under proposed *AAwSOS* (the orange bars in Fig. 7c) the ASPs are significantly less. It can be observed that *AAwSOS* suggested by our *Hide-n-Seek* algorithm showcases higher ASPs as the bus load increases. It is evident from the experiments that the modification in ECU level task schedule by *Hide-n-Seek* manages to reduce ASPs by a higher margin when compared with attack-unaware randomization policies like [11], [16].

V. CONCLUSION

In this work we address an important security vulnerability prevalent in the automotive domain, i.e. the impact of bus-off attack on real time control tasks. We propose a lightweight algorithmic framework *Hide-n-Seek* that applies a set of attack-aware transformations on the task execution schedules of ECUs in order to effectively obfuscate the CAN message schedule. This reduces the bus-off attack success probability and incorporates a method to detect such attempts. We plan to extend this work with thorough evaluation of task scenarios executing under different real time scheduling strategies while benchmarking the attack detectability, i.e. false-alarm and true positive rates.

REFERENCES

[1] S. Hounsinou *et al.*, "Vulnerability of controller area network to schedule-based attacks," in *RTSS*, IEEE, 2021.

[2] S. F. Wiesbaden, "AUTOSAR — the worldwide automotive standard for e/e systems," *ATZextra worldwide*, vol. 18, pp. 5–12, Oct 2013.

[3] K. Koscher *et al.*, "Experimental security analysis of a modern automobile," in *S & P*, pp. 447–462, IEEE, 2010.

[4] K.-T. Cho *et al.*, "Error handling of in-vehicle networks makes them vulnerable," in *CCS*, 2016.

[5] K.-T. Cho and K. G. Shin, "Viden: Attacker identification on in-vehicle networks," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1109–1123, 2017.

[6] K.-T. Cho *et al.*, "Fingerprinting electronic control units for vehicle intrusion detection," in *USENIX Security)*, 2016.

[7] S. Kulandaivel, T. Goyal, A. K. Agrawal, and V. Sekar, "{CANvas}: Fast and inexpensive automotive network mapping," in *28th USENIX Security Symposium (USENIX Security 19)*, pp. 389–405, 2019.

[8] X. Ying, G. Bernieri, M. Conti, and R. Poovendran, "Tacan: Transmitter authentication through covert channels in controller area networks," in *Proceedings of the 10th ACM/IEEE International Conference on Cyber-Physical Systems*, pp. 23–34, 2019.

[9] H. J. Jo, J. H. Kim, H.-Y. Choi, W. Choi, D. H. Lee, and I. Lee, "Mauth-can: Masquerade-attack-proof authentication for in-vehicle networks," *IEEE transactions on vehicular technology*, vol. 69, no. 2, pp. 2204–2218, 2019.

[10] K. Serag, R. Bhatia, A. Faqih, M. O. Ozmen, V. Kumar, Z. B. Celik, and D. Xu, "Zbcan: A zero-byte can defense system," pp. –, 2023.

[11] K. Krüger *et al.*, "Vulnerability analysis and mitigation of directed timing inference based attacks on time-triggered systems," *LIPIcs*.

[12] N. Vreman, R. Pates, K. Krüger, G. Fohler, and M. Maggio, "Minimizing side-channel attack vulnerability via schedule randomization," in *2019 IEEE 58th Conference on decision and control (CDC)*, pp. 2928–2933, IEEE, 2019.

[13] M. Nasri *et al.*, "On the pitfalls and vulnerabilities of schedule randomization against schedule-based attacks," in *RTAS*, IEEE, 2019.

[14] S. Ghosh *et al.*, "A structured methodology for pattern based adaptive scheduling in embedded control," *ACM TECS*, 2017.

[15] K. Serag, R. Bhatia, V. Kumar, Z. B. Celik, and D. Xu, "Exposing new vulnerabilities of error handling mechanism in can.," in *USENIX Security Symposium*, pp. 4241–4258, 2021.

[16] M.-K. Yoon, S. Mohan, C.-Y. Chen, and L. Sha, "Taskshuffler: A schedule randomization protocol for obfuscation against timing inference attacks in real-time systems," in *2016 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pp. 1–12, IEEE, 2016.