

DYST (*Did You See That?*): An Amplified Covert Channel That Points To Previously Seen Data

Steffen Wendzel , Member, IEEE, Tobias Schmidbauer , Sebastian Zillien , and Jörg Keller 

Abstract—Covert channels are stealthy communication channels that enable manifold adversary and legitimate scenarios, ranging from stealthy malware communications to the exchange of confidential information by journalists. We present DYST, which represents a new class of covert channels we call *history covert channels* jointly with the new paradigm of covert channel *amplification*. All covert channels described until now need to craft seemingly legitimate flows or need to modify third-party flows, mimicking unsuspecting behavior. In contrast, history covert channels can communicate by *pointing to unaltered legitimate* traffic created by regular network nodes. Only a negligible fraction of the covert communication process requires the transfer of actual covert channel information by the covert channel’s sender. This information can be sent through different protocols/channels. Our methodology allows an *amplification* of the covert channel’s message size, i.e., minimizing the fraction of *actually transferred* secret data by a covert channel’s sender in relation to the *overall* secret data being exchanged. Further, we extend the current taxonomy for covert channels to show how history channels can be categorized. We describe multiple scenarios in which history covert channels can be realized, analyze the characteristics of these channels, and show how their configuration can be optimized.

Index Terms—Covert channel, steganography, information hiding, network security, internet censorship, amplification attacks, anomaly detection.

I. INTRODUCTION

COVERT channels are policy-breaking and stealthy communication channels that are not foreseen in a system’s design [1], [2], [3]. Such channels are regularly used to transfer secret information, e.g., for the purpose of data exfiltration or malware communications [4], [5], [6], [7]. However, covert

channels can also be applied for censorship circumvention, e.g., by journalists [4], [8]. Covert channels have been investigated for different environments, including networks [4], [9], [10], cyber-physical systems [11], [12], [13], [14], local processes/systems [15], [16], [17] and in out-of-band scenarios, such as sound, light, vibration, radio-frequency, magnetic fields or temperature [5], [18], [19], [20], [21], [22]. Similarly to covert channels, some physical-layer security (PLS) methods also aim at providing confidential communication, e.g., through introduction of friendly jamming/noise [23], [24]. Covert channel signals also aim to hide within noise. The core difference to PLS is that an adversary is explicitly assumed to be aware of the artificial noise while a covert channel aims to prevent that the *presence* of a secret communication is perceptible to the adversary, at all.

All known covert channels *embed* secret messages into flows. Since their first appearance in the 1970 s [1], authors have performed one of two embedding actions: 1) they either relied on the creation of *own* traffic (so-called *active* sending), into which they embed the secret data; or 2) they *modified* legitimate traffic (or its characteristics) transmitted by third-party nodes to embed the secret data (called *passive* sending). In both cases, the embedding of secret information renders a channel slightly detectable. This aspect is the central limitation of all previously known covert channels.

To amend this limitation, it would need a covert channel that (at least partially) transfers secret data through *unmodified* legitimate traffic generated by third-party nodes. In other words, it would require that the covert channel’s sender would craft only a small fraction of own traffic (or modify only a small fraction of third-party traffic) while taking advantage of traffic it neither generates nor modifies.

In this paper, we present a new covert channel called DYST (*Did You See That?*) that fulfills this criterion for the first time since Lampson founded this research area in 1973. In particular, our contributions are as follows:

- 1) *Novel Class of Covert Channels and Concept of Covert Channel Amplification*: We introduce *history covert channels*. These covert channels advance over state-of-the-art as they do not modify existing traffic and generate only minimal covert channel traffic. This traffic is solely used for informing (*signaling*) a receiver that a secret message *appears*, while *no* own data traffic (represents the *actual* secret message) is transferred. This is used to *amplify* the size of the secret message, i.e., the signal is smaller than the actual secret message and thus more challenging to detect.

Manuscript received 22 July 2023; revised 29 March 2024; accepted 30 May 2024. Date of publication 6 June 2024; date of current version 16 January 2025. The work of Steffen Wendzel and Sebastian Zillien was supported by a joint initiative of the Federal Government and the German States through “Innovative University” programme, under Project EMPOWER, FKZ 03IHS242D. This work was supported by the HS Worms library funds. (*Corresponding author: Steffen Wendzel.*)

Steffen Wendzel is with the Center for Technology and Transfer (ZTT), Hochschule Worms, 67549 Worms, Germany, and also with the Faculty of Mathematics & Computer Science, FernUniversität, 58097 Hagen, Germany (e-mail: wendzel@hs-worms.de).

Tobias Schmidbauer is with the Faculty of Computer Science, Nuremberg Institute of Technology, 90489 Nürnberg, Germany (e-mail: tobias.schmidbauer@th-nuernberg.de).

Sebastian Zillien is with the Center for Technology and Transfer (ZTT), Hochschule Worms, 67549 Worms, Germany (e-mail: szillien@hs-worms.de).

Jörg Keller is with the Faculty of Mathematics & Computer Science, FernUniversität, 58097 Hagen, Germany (e-mail: joerg.keller@fernuni-hagen.de).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TDSC.2024.3410679>, provided by the authors.

Digital Object Identifier 10.1109/TDSC.2024.3410679

- 2) *Scenario Provision*: We provide multiple scenarios for local network and remote communications to show where the application of history covert channels can be beneficial. In particular, we show that history covert channels can be applied in highly constrained environments, where a covert sender is incapable of creating any new/modifying any existing traffic and is thus only allowed to send packets with pre-defined content. These features render history channels challenging to detect.
- 3) *Taxonomy Extension*: We extend the existing taxonomy for network covert channels with a new category called *fully-passive sending* to reflect all components of this new class of covert channels.
- 4) *Functionality Description for Multiple Methods*: As history covert channels comprise a whole family of variants, we provide a description of their functionality using different variants called DYST-Basic (for local networks), DYST-Ext (also for local networks), DYST-Remote-Smarthome (for local-to-remote connections) and DYST-Remote-RTCP (for local-to-remote connections).
- 5) *Theoretical Analysis*: We conduct a theoretical analysis of the performance and optimization of DYST.
- 6) *Implementation*: We describe several ways to implement a history covert channel in a network environment. As an example, we provide the first implementations of such a channel for DYST-Basic and DYST-Ext. Our implementations contain a *data* channel that requires no own or modified traffic and a *signal* channel that consists solely of rarely sent packets with legitimate content (representing only 1 covert bit).
- 7) *Evaluation*: Using both local network implementations, we evaluate DYST’s robustness, detectability, and optimization under different settings and show that DYST-Basic and DYST-Ext allow the transfer of variable secret data bits through only 1 covert signaling bit, which state-of-the-art covert channels do not achieve. We further simulate a remote version of DYST-Basic called DYST-Remote-Smarthome to show the feasibility of history channels outside of local networks as well as a throughput-enhanced multi-pointer variant.

The remainder of this paper is structured as follows. Section II provides background information and discusses related work while Section III presents the functioning as well as the theoretical description and optimization of DYST. We describe our experimental testbed and the evaluation of DYST in local networks in Section IV. Next, we conduct a feasibility-analysis of a remote scenario for DYST in Section V and discuss a performance optimization in Section VI. A discussion is provided through Section VII. Finally, Section VIII concludes. The electronic supplement covers additional optimization aspects of DYST.

II. BACKGROUND & RELATED WORK

Covert Channels: A *Covert Channel* exchanges information in a stealthy manner between a *covert sender* (CS) and one or more *covert receiver(s)* (CR). A covert channel is one that is

not foreseen in a system’s design [1] and relies on the concept of policy-breaking communication [2]. If the covert information can be received by more than one CR, the communication can be considered as a multicast or even broadcast covert channel. In computer networks, the covert channel nests into a network protocol, e.g., by manipulating bits of a packet header or by adjusting the delays between successive network packets. Similar to covert channels, some physical-layer security methods utilize nearby network nodes’ emitted (artificial) noise [23], and steganography tools typically aim to hide within noise, which also applies to DYST’s signaling channel (but not its data channel).

In [25], [26], [27], Cabuk et al. propose the idea of an advanced version of a covert channel based on delays between packets: the covert channel transmits the hidden data by modulating the delays between consecutive network packets. The advanced version of their covert channel mixes covert transmissions with sections of real, legitimate network traffic. This helps to skew the statistics and makes the detection of the covert channel harder. The difference to our approach is that Cabuk et al. use sections of legitimate traffic solely to introduce noise into the actual covert channel signal. The legitimate traffic carries no hidden information at all.

Several additional methods work similarly to the one of Cabuk et al. For instance, *JitterBug* by Shah and Molina [28] adds random delays to legitimate Telnet traffic. Walls et al. proposed *Liquid* [29], an extension of JitterBug, in which they split the channel into “transmitting” and “shaping” delays (shaping delays carry no information but manipulate the statistics of traffic). Similarly, Gianvecchio et al. [30] tailor traffic automatically based on the statistical characteristics of legitimate traffic. In all these cases, artificial modifications are performed to transfer secret information, even if based on legitimate traffic, which is the key difference to DYST.

There are also approaches that work on a more abstract level. Yarochkin et al. [31] proposed the so-called *network environment learning* phase. This approach was used solely to determine which protocols occur regularly in a network to succeedingly exploit only these protocols for covert communication. No work is known that exploits legitimate traffic for a covert channel. Moreover, the covert channel of Yarochkin et al. did not split the covert channel’s control channel from its data channel as we do.

Image steganography uses a variant called *cover selection*, where a database of images is used, a hash function is applied to each image, and if the secret message matches the hash value, the image is sent by the covert sender [32]. Several methods have been proposed to conduct or optimize the cover selection process, e.g., [33], [34], [35], [36], [37], [38]. Similarly, *coverless image steganography* utilizes a database of image patches to (partially) reconstruct an original secret message using image patches, see, e.g., [39] and [40]. In contrast to cover selection and coverless image steganography, our method operates in a network and it relies on network traffic *that is transmitted anyway*, i.e., independent of the steganographers, and only uses signals, e.g., ARP or RTCP requests, instead of generating extra traffic such as images, even if those images are innocent.

The Address Resolution Protocol (ARP) is used to associate a link layer address, e.g., MAC (Medium Access Control) address to an IPv4 address, and to this end uses request and reply packets on the link layer. ARP had been utilized in [41] to implement a local covert channel by encoding covert information into the target IP field of an ARP request, sending the request directly to the covert receiver. Another approach utilizing ARP for covert communication had been described in [42], wherein ARP requests are exploited to store covert information within ARP tables of uninvolved intermediary third-party nodes. The covert information have to be fetched by the covert receiver utilizing SNMP. The Real-Time Protocol (RTP) and the Real-Time Control Protocol (RTCP) are used for media streaming. We use these protocols for our PoC remote scenario. These protocols had been exploited in [43], [44], [45] for covert communication. In [43], the authors utilized various ways to directly send covert information to a covert receiver by crafting and manipulating RTP traffic. The techniques include beside others the manipulation of timing behavior, and embedding information in the padding, the extension and the sequence number field. In [44], the jitter field is manipulated for directly sending information to the covert receiver, and in [45], the authors significantly improve RTP and RTCP based timing channels. While all these covert channels had been implemented with ARP or RTP/RTCP, every former implementation had to *craft or manipulate* packets of these protocols to transport covert information to covert receivers. The approach described in *our* paper significant differs from these former implementations as there is no need to craft a packet carrying the covert data to transmit to the covert receiver, as legitimate packets carry the covert, but still legitimate information that are solely signalled to be part of a secret message.

There is one method that actually splits the control channel from the data channel: as shown by Wendzel and Keller in [46], several covert channels propose to utilize internal control protocols. Therefore, a covert channel is nested into the utilizable bit areas of a network packet. Some of the utilizable bits are used for the control protocol while others are used for the data channel. However, that approach has a major limitation compared to ours as both, the control and the data channel, reside in the same packet and modify legitimate packets or craft new packets instead of exploiting solely unaltered legitimate traffic for the data channel.

Finally, there is one proposal by Caviglione et al. from botnet research that works by waiting for a pre-defined network packet sequence [47]. If the sequence occurs, the botnet nodes would perform a certain action. The idea was solely described on a conceptual level and was not implemented by the authors. Further, their concept did not involve the option to influence which secret message is transferred, which makes it fundamentally different from our history channels.

HICCUPS [48] points to data *within* a network packet. Therefore, a sender corrupts a frame checksum, which indicates the presence of a secret message in the frame's payload. HICCUPS is tied to LAN environments and the frame-level. It further does not point to historic data but solely *places* the pointer *and* the secret data inside the packet (instead of finding matching

third-party data). Thus, HICCUPS is not a *history* covert channel and provides no amplification.

Network Covert Channel Detection: Several detection methods for covert channels have been proposed throughout the years. Popular ones are, e.g., compressibility score [26], ϵ -similarity [27], regularity metric [27], a method from Berk et al. [49], as well as classical methods, such as Kullback-Leibler divergence test [50], Kolmogorov-Smirnov [50] or entropy-based analyses [10]. All of these methods require at least a few hundred covert channel packets to provide somehow reliable detection of covert channel flows. In contrast, history channels send few *signaling* packets per time, resulting in only a minimal influence on a flow while the data flow of history channels is entirely legitimate and thus indistinguishable. We evaluate two common detection methods on DYST in Section IV.

III. THE HISTORY COVERT CHANNEL METHOD

In this section, we first discuss the requirements of our history covert channel, followed by a description of the detailed functionality of DYST. Further, we explain the chosen parameters for DYST, including the optimization, and finally extend the existing taxonomy of active and passive covert channels.

Definition 1: A *history covert channel* is one that *points* to already existing (live or stored) data that matches a secret message instead of *sending* a secret message itself. The only covertly transferred information is the pointer. ■

Data, such as packets, that do *not* match a secret message, are not processed further by a covert sender, i.e., no pointer will refer to these. Note that pointers may be considered as a modification of third-party or a crafting of own traffic. However, the content that is pointed to is legitimate. Sending small-bit pointers to secret data with *more* bits thus minimizes the fraction of covert traffic to be sent/modified in comparison to covert data transferred overall, which we believe is a key novelty in this domain. This means that history covert channels perform a form of covert message size *amplification*. We call the size-increasing factor for the covert message the *covert amplification factor* (CAF). The functioning of a history covert channels that uses legitimate broadcasts as a data channel is visualized in Fig. 1. Note that non-broadcast data channels can be used too, which we will detail for remote scenarios in Section III-A.

Note: We chose the specific *proof-of-concept* (PoC) implementation discussed in the remainder to simplify experiments and explanations. We note, however, that history covert channels are not restricted to local networks and more variants are possible, cf. Section VII.

A. Threat Model and Requirements

In general, we assume a scenario where it is beneficial for a CS to send as few secret information bits to a CR as feasible. In particular, we consider a scenario where only legitimate-*content* messages are directly exchanged between CS and CR and that only the *timing* of certain messages is used as a referrer to some other message containing the actual secret content. Thus, we assume a hybrid covert channel, split into a signaling and a data

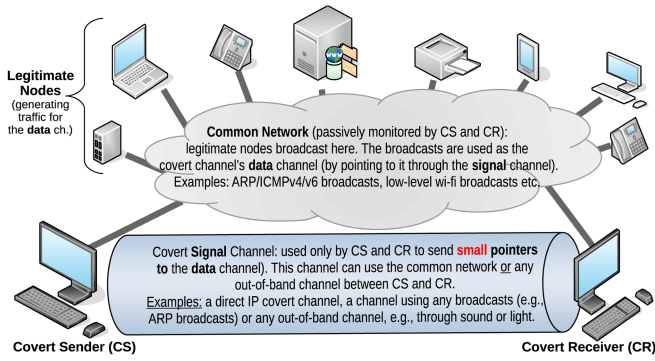



Fig. 1. General functioning of a history covert channel that uses broadcasting. The amplification is achieved by sending small pointers that refer to larger data pieces (e.g., packets or their hash values) in the data channel.

LAN Scenario

Step 1: Legitimate traffic is broadcasted



Step 2: CS refers CR to the previously seen broadcast through some other legitimate broadcast's timing (e.g., a typical ARP request)

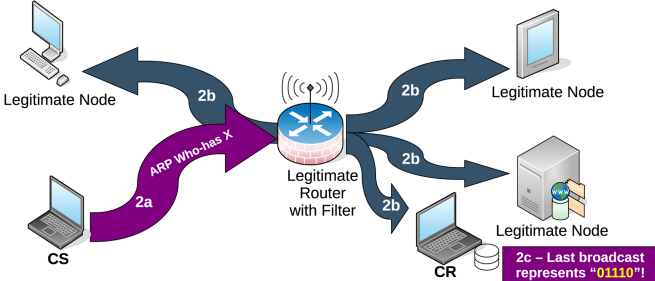


Fig. 2. Confined communication in LAN: CS is not capable of placing the secret message content by itself due to restrictions.

channel. Our channel does not match any of the recent patterns of indirect covert channels found in [51].

Several scenarios are imaginable where using such a covert channel would be beneficial:

Local scenario 1/1 (DYST-Basic and DYST-Ext): Highly constraint LAN environments, where a CS cannot freely craft the content of network packets at will due to a filter and can barely manipulate the timing of some legitimate messages. This scenario is shown in Fig. 2. CS could try to establish a covert timing channel with CR in such a scenario. However, typical covert timing channel detection methods might detect a classical timing channel. For this reason, DYST’s timing channel does not contain the actual secret information. Instead, it is bound to the occurrence of legitimate messages called *packets of interest*

(*PoI*) exchanged within the network environment and visible to CS and CR: CS only sends a timing signal if some legitimate node transfers a message whose parts matches the content of a secret to be transferred and that is also received by CR, such as a broadcast (step 1a in Fig. 2). For this reason, the timing channel messages (signal messages) of CS only occur after such legitimate messages appear, and thus do not reflect a typical timing channel pattern. In step 1b, the legitimate message (of step 1a) is broadcast to all other nodes. The CR keeps track of all recently seen broadcasts in a local database (packets older than a few seconds can be discarded). Since the broadcast message’s content matches the secret message to be transferred, CS sends an own broadcast (step 2a), which is also broadcast to all nodes of the LAN, including CR (step 2b). In step 2c, CR loads the most recent legitimate broadcast (that was *not* sent by CS) from its database and extracts the secret information from that packet. Note that due to the secret message being represented by some third-party traffic, CS does not need the capability to craft such messages by itself since it *points* to these messages instead using the signal.


Remote Scenarios Second, there are non-local scenarios imaginable, where DYST-Basic would be applicable but utilizes different protocols.

Remote scenario 1/2: Highly-constrained remote access/censorship circumvention (DYST-Remote-Smarthome). Similarly to the local scenario, we assume that CS operates in a restricted environment, where only pre-defined packets can be sent by CS and where CS can influence the time at which these are sent. For instance, some smart devices regularly communicate from internal networks to public services, such as cloud providers. This is a *Man-in-the-Middle* (MitM) setup, where an on-path CS would be located in the LAN (alternatively on a router or on the path to the destination) and where CR would be located down-path outside the LAN, e.g., in the public Internet, close to the destination or residing directly at the destination. In other words, CR must have direct access to the traffic sent by the legitimate sender. This scenario is shown in Fig. 3. Note that no broadcast traffic could be used in this scenario as it would not reach the CR. Therefore, CS needs to wait for regular unidirectional (or multicast) messages (step 1a) from the legitimate sender to a destination where the traffic passes CR or where CR is the destination. As in the previous scenario, CR records all recent *packets of interest* (PoI) in a local database (step 1b). Afterwards, CS needs to send some allowed packet observable by CR as a signaling packet (step 2a), such as a whitelisted (i.e., not filtered) DNS or HTTP request, as a *signal* to CR (instead of a broadcast message in the previous scenario). For instance, if CR is a web service that has the current news or the current canteen menu on its webpages, then CS could time the specific allowed requested to CR.¹ Alternatively, if CS is not allowed to send packets to CR directly, CS might send data to a third-party node observed by CR. For instance, CS and CR could use a network entity, e.g., a feed where CS is allowed to “like”

¹We assume that HTTP requests face deep packet inspection (DPI) or are normalized so that CS has indeed no option to modify pre-defined requests of a set of allowed network services CS can use.

LAN-to-Remote Site Scenario

Step 1: Legitimate traffic that can be observed by both, CS and CR, is sent by a legitimate sender



Step 2: CS refers CR to the previously seen packet through some legitimate signal message.




Fig. 3. Confined communication from a LAN to a remote site: CS resides inside the local network and CR resides in an uncensored public network.

postings (but not posting any content) so that the timestamp of a like is observable by CR to establish the signal channel. Finally, CR interprets the last observed packet's secret content (step 2b) to extract the secret message.

Remote scenario 2/2: Censorship circumvention using RTCP (DYST-Remote-RTCP). This scenario is a special form of *remote scenario 1* and added solely to underpin the existence of this option. This scenario is called DYST-Remote-RTCP and can be imagined in settings where sender and receiver can establish a communication which includes at least a partial time synchronization, e.g., a media communication via RTP in conjunction with RTCP [52]. The PoI comprise the media traffic itself which is transported via RTP. A time synchronization is achieved via Sender Reports (SR) in RTCP which are regular messages that include a Network-Time Protocol (NTP) timestamp and counters for how much traffic has been sent in the last period. Pointers are signalled within such SR. CS and CR listen to traffic on both protocols, and CR records from RTP (which hosts the majority of data traffic) all PoI of the time interval since the last SR. CS only needs to signal within such a SR if a suitable piece of information has been sent within a PoI since the last SR. The pointer to this information may consist of an offset, or may be a part of the information itself that is sufficient to uniquely identify the information within all PoI since the last SR. Thus, the history channel at least reduces the fraction of secret bits that the covert sender needs to transfer explicitly through such a covert channel in contrast to a covert channel that uses RTCP alone (cf., e.g., [53]), and still has no need to modify traffic in RTP protocol itself.

Environment Characteristics: DYST relies on some crucial characteristics of the environment in which CS and CR operate:

- 1) CS and CR must be able to observe some messages which they both receive (almost) at the same time. There are different options to achieve that: i) CS and CR could read broadcast messages in a local WiFi network that both of them receive; ii) CS and CR could only evaluate messages that pass through the routing path of CS and CR (e.g., when both act as routers); iii) CS and CR are members of a multicast-group, receiving frequent updates (for instance IGMP, like exploited for an active indirect covert channel in [54]). In any case, CR must be able to store information about the recently received packets in a local database (older packets can be discarded from the database);
- 2) To assign the same secret data to observed packets, CS and CR need to utilize the same procedure. This means that they need to read the same input parameters (i.e., which packets and which header and body fields of these packets, timestamps, and so forth are used) of the network traffic and also apply the same hashing function over these data (the hash can then represent secret data);
- 3) Moreover, CS must be able to send the signal so that it is observable by CR in a timely manner, e.g., CS might be able to send a legitimate ARP (or, e.g., DHCP, ICMPv6) broadcast to the local WiFi network or a remote message. While it is not within our scope, CS might alternatively use some an out-of-band covert channel [5] to signal CR.

Optionally, CS and CR may even reside on two existing legitimate systems as there is no need for CS and CR to utilize added systems. The above-mentioned conditions can also be fulfilled outside the scope of networks.

Adversary Capabilities: We assume that the adversary (e.g., a government-level Internet censor) has highly restrictive capabilities:

- 1) The adversary is able to *observe all* traffic exchanged in the local network, where CS resides.
- 2) The adversary can further *block any* traffic sent by CS if its content (that is a combination of packet header and payload) does not match an element of a set of pre-defined legitimate messages. For instance, in a local network scenario, CS might be able to only send one specific ARP broadcast that calls for the address of the local network's router.
- 3) The adversary is also able to observe *all* traffic that is exchange on the path between CS and CR, even if CR resides in a remote network. If the adversary detects a suspicious communication, it can decide to block the communication from CS to CR for an arbitrary time.
- 4) The adversary is not capable of locally monitoring actions or resources on CS or CR, i.e., there is no adversary malware on CS or CR.

B. Functioning of DYST

In this section, we primarily describe two essential variants of DYST: first, we explain DYST-Basic, a simple variant where all bits of an observed hash need to match the chunk of the

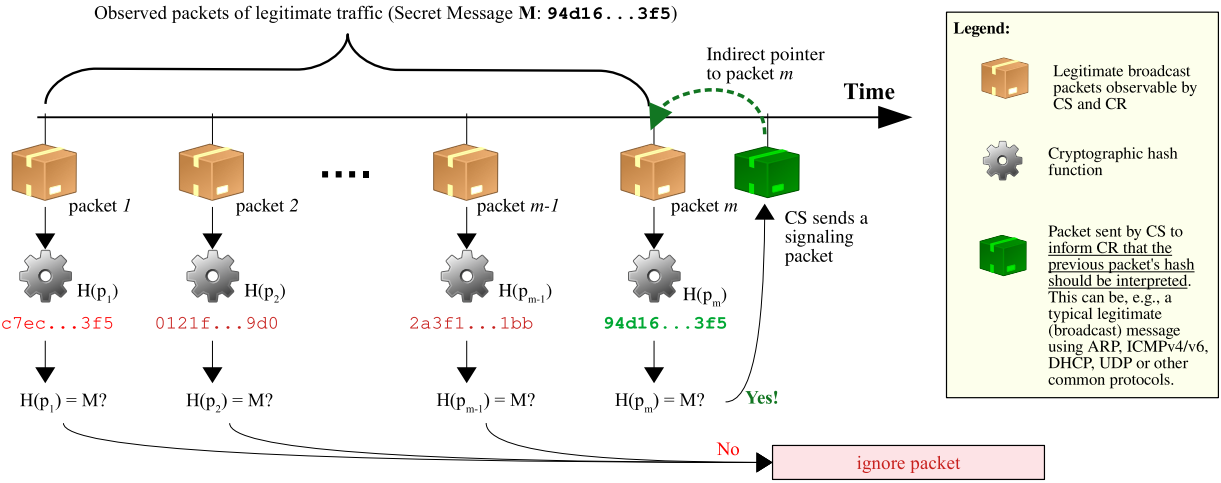


Fig. 4. The DYST-Basic Sending Process. Hashes representing a secret message can be calculated over whole packets or just parts of packets (e.g., selected header fields). Note that packets with non-matching hash values are not processed further. Signaling packets can take the form of any typical (legitimate) broadcast, such as ARP or DHCP requests.

TABLE I
TABLE OF USED NOTATIONS IN THIS PAPER

Symbol	Definition
$ \cdot $	length of a string . or cardinality of a set .
bw_{basic}	bandwidth of DYST-Basic
bw_{ext}	bandwidth of DYST-Ext
$C(\cdot)$	checksum function on input .
c	number of bits used for error correction
D	sender-side robustness timeframe between two packets of interest (PoI)
$d(x, y)$	hamming distance between bitvectors x and y
$dist_{basic}()$	signaling distance of DYST-Basic
$dist_{ext}()$	signaling distance of DYST-Ext
$extract_{sm}(x)$	function to extract the last c bits from a bitvector
$extract_{msg}(x)$	function to extract the first $h - c$ bits from a bitvector
$F(a)$	cumulative distribution function of a random variable X , i.e., $P(X \leq a)$
$H(\cdot)$	cryptographic hash function using input .
h	length of hash value
h_i	hash value i
\mathfrak{S}	compressing function
κ	compressibility score
k	number of message M 's fragments
$len(M)$	length of message M
M	secret message
\tilde{M}	encoded message M with concatenated checksum
$P(X)$	probability function for X
p_i	network packet i
R	receiver-side robustness timeframe between two PoI
S	string representation of concatenated inter-packet delays
S_{h_i}	set of bit vectors with distance $\leq t$ to h_i
$T_{h,t}$	number of possible modifications of a hash value to test
t	number of non-matching bits of a message
$U_{h,t,c}$	chance that message M has the first fit among $T_{h,t}$ possibilities
X	random variable
x	bit vector of length h

covert message. Second, we will describe DYST-Ext, which is an extended variant capable of transferring correct information, though the hashes are not matching perfectly. This approach creates more variants to choose from in the tradeoff between steganographic bandwidth and detectability. Further, we describe derivatives of DYST-Basic: DYST-Remote-Smarthome and DYST-Remote-RTCP for the remote scenarios 1 and 2 (matching Fig. 3). Note that an overview of the notations used in this paper is given by Table I. Further, Table II summarizes how each DYST variant is handled in this paper.

1) *DYST-Basic*: For DYST-Basic, the following steps must be performed continuously until a whole message is transferred (see Fig. 4):

- 1) CS and CR both record legitimate traffic with their network interfaces connected to a shared network.
- 2) For each packet of interest p_i that both CS and CR can observe (e.g., broadcast messages), they apply a hash function $H(\cdot)$ to the input values they have agreed on, to generate a hash value $h_i = H(p_i)$ of length h .

CS and CR can exchange a secret message in a bi-directional manner. In the remainder, we describe the sending process from CS to CR. However, CR can simultaneously operate as a CS and CS as a CR to send/receive data.

The sending process can be described as follows (see Fig. 4) and requires that each secret message M to be sent has the length h . To transfer a secret message of length $k \cdot h$, $k \in \mathbb{N}$, $k \geq 2$, the message must be split into k fragments and the sending process must be performed successively for each of the k fragments.

- 1) To signal a secret message of length h , CS waits for a packet of interest p_i with a hash value $h_i = H(p_i)$ which equals M .
- 2) After the CS observed such a packet, it sends a signal, for example a legitimate ARP request in which it asks for the address of some legitimate node (e.g., a router) in the network. Other possibilities for signals are discussed in Sections III-B3 and III-B4.

Finally, the receiving process is conducted as follows:

- 1) CR interprets the occurrence of the signal message from CS as a prompt that the expected message can be observed in the data channel.
- 2) CR interprets the hash value of the previous packet of interest that represents the covert message.

Obviously, the channel is noisy and requires error detecting (and correcting) bits or mechanisms to ensure a robust transmission of the correct information.

The major advantage of such a history covert channel's sending procedure is that CS needs to send only one bit of covert

TABLE II
OVERVIEW OF DYST VARIANTS IN THIS PAPER

DYST variant	Matching Scenario (Section III-A)	Functionality Description (Section III-B)	Optimization (Section III-C)	Evaluation Method for Covert Channel's Characteristics (Section IV / Section V)
Basic	Constrained LAN	✓	✓	Implementation, with partial simulation
Ext	Constrained LAN	✓	✓	Implementation, with partial simulation
Remote	Constrained Network with Remote Link / Censorship Circumvention	✓ (essentially the same as DYST-Basic)		Due to overlap with DYST-Basic: only simulation with testbed traffic from a smart home

information (represented by the existence of the signal message, e.g., an ARP request or RTCP packet) to transfer h bits of secret content, i.e., the CAF for the covert message's size is h (or a fraction of h if more than one signaling bit must be sent). DYST modifies the inter-packet gaps by introducing signal messages, while existing timing channels modify timing of existing packets to transmit a bit of the secret message. We will treat this in more detail when analyzing detectability, cf. Section IV-E.

DYST's example ARP request can be replaced by any other seemingly legitimate unicast, broadcast, multicast, or anycast message, observable by the CR.

Moreover, CS reaches multiple CR simultaneously, if desired, rendering the channel a multicast or broadcast covert channel.

Assuming that each of the 2^h possible hash values is equally likely, then on average an exact match between secret message M and hash value h_i will be achieved after 2^h packets. If an exact match is not required but up to t bits can be wrong, i.e., the hamming distance between secret message M and hash value h_i can be at most t , then the chance of such a partial match would increase by a factor $\sum_{i=0}^t \binom{h}{i}$. To enable the CR to still decode the message correctly, the message would have to be encoded with an error-correction code that allows correction of up to t bit errors. This means that only $h - c$ bits are available for the secret message itself and the remaining c bits are used for the error-correction bits [55]. For a binary block code, we have $c > 2t$ as $c = 2t$ is a bound achieved by *maximum distance separable* (MDS) codes, which however only exist in trivial form for binary² codes [56, Prop. 9.2]. Furthermore, besides a strong restriction on the number of message bits (the larger c , the smaller $h - c$), only some combinations of h , c , and t are available for applicable code families such as binary *Bose-Chaudhuri-Hocquenghem* (BCH) code, see e.g., [55, App. A]. Our initial investigations revealed that error-correction codes only in some cases match the performance (in terms of bandwidth and average signal distance) of DYST-Basic, and mostly perform worse. Yet we carried over the idea of using partial matches to using checksums instead of error-correction codes, creating an extended version of DYST.

2) *DYST-Ext*: The functionality of DYST-Ext works similarly to DYST-Basic, but the secret message M now only comprises $h - c$ bits (this reduces the CAF by c bits in comparison to DYST-Basic), and is concatenated by CS with a c -bit checksum to an encoded message \tilde{M} of length h . When comparing \tilde{M} with h_i , CS allows up to t non-matching bits. Thus, the advantage of DYST-Ext comes from the fact that it can utilize a larger fraction of observed messages, leading to a shorter waiting time

²Using non-binary codes decreases match probabilities further and thus is no option, either.

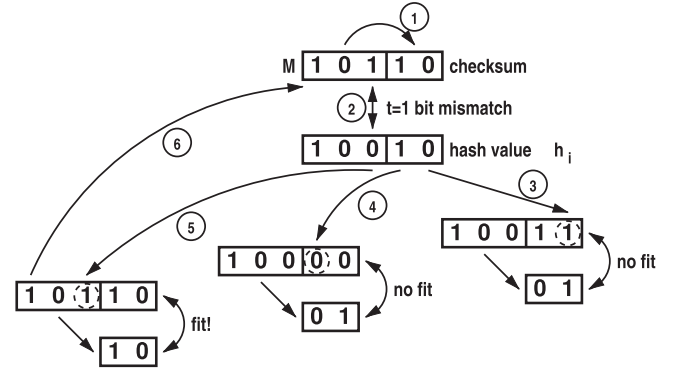


Fig. 5. Functioning of DYST-Ext. For illustration, we use a message chunk M of $h - c = 3$ bit length and a checksum of $c = 2$ bit, which simply represents the number of ones in the message chunk as a binary number (step 1). In the hash value h_i with $h = 5$ bits, only $t = 1$ bit do not match (step 2). CS thus checks all modifications of the hash value until a hit between the reconstructed message and reconstructed checksum occurs (steps 3 to 5, flipped bits marked by dashed circles, order of bit flips: right to left). As the reconstructed message equals message chunk M (step 6), CS will send an ARP request to CR. CR will perform the same computation and reconstructs M . For message chunk $M = 110$ (same checksum), the same fit would apply but not reconstruct M , and no ARP is sent.

for fitting packets. CR, upon receiving a signal (such as an ARP request) indicating a secret message transfer, again picks up the latest hash value h_i from the hash database. It then tries out all possibilities to flip up to t bits in h_i , until the checksum of the first $h - c$ bits in the modified h_i matches the last c bits in modified h_i (called a *hit*) for the first time. CR accepts the first $h - c$ bits as secret message M .

CS knows the order in which CR will apply modifications to the hash value h_i . Thus, CS can check if the first hit really will produce the message M . CS will only send an ARP request as a signal for CR if at most t bits of h_i and \tilde{M} do not match and the first hit found by CR will produce M . So, not all t bits matching packets can be utilized for this approach.

Fig. 5 illustrates the working of DYST-Ext.

Formally, if $C : \{0, 1\}^* \rightarrow \{0, 1\}^c$ is the checksum function, and $d(x, y)$ is the hamming distance between two bitvectors x and y of equal length, CS first checks if

$$d(M || C(M), h_i) \leq t. \quad (1)$$

If so, CS enumerates the set

$$S_{h_i} = \{x \in \{0, 1\}^h \mid d(x, h_i) \leq t\}, \quad (2)$$

in a pre-defined order, i.e., it generates a sequence of distinct bitvectors $x^{(1)}, x^{(2)}, \dots$ that together form S_{h_i} . Let $extmsg(x)$ be a function to extract the first $h - c$ bits from a bitvector of

length h , while $extchksm(x)$ extracts the last c bits. For $j = 1, 2, \dots, |S_{h_i}|$, CS checks if

$$C\left(extmsg\left(x^{(j)}\right)\right) = extchksm\left(x^{(j)}\right), \quad (3)$$

and stops with the first hit at index j^* . If

$$M = extmsg\left(x^{(j^*)}\right), \quad (4)$$

then CS sends an ARP request.

Upon receiving such an ARP request, CR looks up h_i , and also enumerates S_{h_i} according to (2), does the computations from (3) and stores secret message M according to (4).

Please note that both DYST-Basic and DYST-Ext are *families* of variants, because they are parameterized in H and h, c, t , respectively.

3) *DYST-Remote-Smarthome*: In case of DYST-Remote-Smarthome, the same considerations apply as for DYST-Basic with the exception that no local broadcasts are considered as PoIs and no ARP broadcasts are sent as signal packets. Instead, PoIs are solely such messages that can be observed by both, CS and CR, i.e., both must be on the path of a third-party packet. Further, signal packets can be any legitimate packets generated by CS that pass CR directly (e.g., if CR is a router) or that can be inferred by CR (e.g., because CS pulls a *Git* code repository while CR can monitor such pulls).

4) *DYST-Remote-RTCP*: Finally, we like to highlight the imagined scenario for a censorship circumventing channel that uses RTP/RTCP. This is another approach to transform DYST-Basic into a remote scenario. It can be realized via a media communication, where CS only signals if the first PoI in the time interval since the last SR is the one that serves as the actual secret message. Alternatively, CS might include a short index which of the PoIs in the time interval is the one to be used as actual secret message. Assume for example that the complete communication on RTP is split into pieces of length $2h$ bits each, and the PoIs are comprised of the h -bit hash values of those pieces. For a data rate of r bits per second and an SR interval of t seconds (usually 5 seconds [52]), there are $T = rt/(2h)$ PoIs of h bits each. To uniquely describe the intended PoI, $\log_2 T$ bits are necessary. Thus, to really save bandwidth compared to a direct covert channel of h bits, we require $\log_2 T \ll h$. Given r and t , a suitable value for h can be determined.

C. Parameter Choice

This section explains the optimization of DYST-Basic and -Ext. Since DYST-Remote-Smarthome and DYST-Remote-RTCP are solely deviations of DYST-Basic, their optimization is not detailed separately.

1) *DYST-Basic*: We first consider the situation of DYST-Basic, where CS signals to CR if the hash value of a network packet matches the secret message exactly. If the hash function H has optimal properties, each bit of the hash value has a value of 0 or 1 with probability 50%, respectively. The probability that h_i equals a given secret message M then is 2^{-h} , as all bits of the hash value can be considered independent. As the hash values of the different packets can be considered uniformly distributed and

independent, the number of packets until a match between hash value and secret message occurs follows a geometric distribution with success probability 2^{-h} , i.e., with expectation 2^h .

2) *Pareto-Optimal Variants*: A covert channel such as DYST-Basic can be characterized by two properties: the *distance*, i.e., the average number of observed packets between two signal messages, which will influence detectability, and the steganographic *bandwidth*, i.e., the number of secret message bits transmitted via one signal message in relation to the distance. Thus, DYST-Basic with parameter h is a family of covert channels with distance $dist_{basic}(h) = 2^h$ and bandwidth $bw_{basic}(h) = h/2^h$. The unit of distance is number of packets of interest, while the unit of bandwidth is number of bits of the secret message per packet of interest. By multiplying the bandwidth with the number of packets of interest per hour in a particular scenario, cf. Section IV-B, an absolute bandwidth is obtained in number of bits of the secret message transported per hour. Similarly, distance must be divided by the frequency of PoI to obtain an absolute distance between signal messages measured in hours.

We would like to maximize both: *signaling distances* because there is a threat of detectability when signal messages such as ARP requests occur too often, and *bandwidth* to increase applicability.³ Yet, increasing distance will reduce bandwidth and vice versa.

Hence, to achieve an optimal compromise between the two parameters, we search for a Pareto front, i.e., a set of non-dominated variants.⁴ Alternatively, we impose a constraint on one parameter and search the optimal value of the other parameter, i.e., we cut the Pareto front with a vertical or horizontal line into two halves, and search the point on the front closest to the border (in the “allowable” half). It is obvious that variants of DYST-Basic with different values for h do not dominate each other, as improving one parameter makes the other worse.

3) *DYST-Ext*: In DYST-Ext, only $h - t$ or more bits of the hash value must match the encoded message, comprised of message chunk and checksum. The number of matching bits is a random variable X that is binomially distributed with h trials and success probability 0.5, and thus the probability that at most t of the h bits do not match is

$$P_h(X \geq h - t) = \sum_{j=0}^t \binom{h}{j} / 2^h. \quad (5)$$

CS and CR must try out

$$T_{h,t} = \sum_{j=0}^t \binom{h}{j} = P_h(X \geq h - t) \cdot 2^h \quad (6)$$

possible modifications of the hash value. For each, the chance that the checksum of the message part of such a modified hash value equals the checksum part of the modified hash value, i.e.,

³The third parameter of the covert channel magic triangle, robustness, is considered a fixture in this optimization, as it will be considered before going into this tradeoff, cf. Section IV-D, and will reduce the number of observable packets per time so that it affects all variants in a uniform way.

⁴A covert channel variant A is dominated by variant B if $distance(A) \leq distance(B)$ and $bandwidth(A) \leq bandwidth(B)$.




Fig. 6. Simulation results for DYST-Ext (left), and zooms for DYST-Basic and DYST-Ext, respectively. Axes give distance, i.e., the average number of observed packets between signalled transmission of secret message parts, and steganographic bandwidth, i.e., the number of secret message bits transmitted with one signal in relation to the distance.

the chance that (3) is fulfilled, is 2^{-c} . Thus, the number of trials until a fit will occur is geometrically distributed with parameter 2^{-c} , yet with a limited range of $T_{h,t}$ trials. The chance that the true message has the first fit thus is

$$U_{h,t,c} = \frac{1}{T_{h,t}} \cdot \sum_{k=0}^{T_{h,t}-1} (1 - 2^{-c})^k = \frac{1 - (1 - 2^{-c})^{T_{h,t}}}{T_{h,t} \cdot 2^{-c}}. \quad (7)$$

As both events (at most t non-matching bits, secret message chunk is re-constructed in first fit) can be considered independent, the chance that a message transfer can be signaled by a signal message is their product, $P_h(X \geq h - t) \cdot U_{h,t,c}$. Taking (6) into account, we see that this product is approaching 2^{h-c} , the probability of a signal in DYST-Basic with a message of length $h - c$, yet the additional solutions can still be non-dominated. As the product probability is independent of the particular hash value, the number of packets until a signal occurs is again geometrically distributed, with the expectation

$$\text{dist}_{ext}(h, t, c) = \frac{1}{P_h(X \geq h - t) \cdot U_{h,t,c}}, \quad (8)$$

and the covert channel has a bandwidth

$$\text{bw}_{ext}(h, t, c) = (h - c) \cdot P_h(X \geq h - c) \cdot U_{h,t,c}, \quad (9)$$

as $h - c$ bits of the secret message can be decoded by CR with each signal.

4) *Pareto Front*: We have computed distance and bandwidth for DYST-Basic with $h = 14, \dots, 20$ and for DYST-Ext with $h - c = 14, \dots, 18$, $c = 6, \dots, 10$ and $t = 1, \dots, 5$, both analytically, and supported by simulations of $5 \cdot 10^7$ hash values, where we counted how often CS signals in each variant considered. The hash values in simulations were generated by successively encrypting a 128-bit value with AES and a fixed 128-bit key, starting with value 0 and using the first h bit of the value as hash value. We used three different checksum functions in simulations (we only use first c bits of longer results): SHA-3, CRC8, and a handcrafted function that cuts the encoded message into pieces of length c , adds those pieces as binary numbers, adds 9, and takes the c lowermost bits of the result. All simulations were repeated with a second seed and results were manually compared to exclude the possibility of artifacts, which however did not show.

The values for h, t, c were chosen to allow comparison between different variants in a restricted range, and to illustrate development of distance and bandwidth over the range for a particular parameter.

Fig. 6(left) provides all of the above variants as points in the plane with distance on the x -axis and steganographic bandwidth on the y -axis, as defined in Section III-C2. All points are quite close together, so we do not have a point “cloud” but still a bit “thicker” line. Thus, DYST-Ext extends DYST-Basic in the sense that the user has more choices in the tradeoff between distance (stealthiness) and steganographic bandwidth than with DYST-Basic alone. This is illustrated in Fig. 6(mid) and (right) that depict zooms into $(10^6, 0.1)$ and $(10^5, 0.001)$, respectively. The points that represent variants of DYST-Ext (blue) fill the gaps between the points representing variants of DYST-Basic (red). Fig. 6(right) depicts the region of interest, i.e., the region where the actual tradeoff between distance and bandwidth can be seen. Put otherwise, this is the region with the bend in the bottom-left quadrant of Fig. 6(mid).

The Pareto front contains only about one-third of the variants. As its shape is similar to Fig. 6(left), we refrain from showing another figure. Among the DYST-Ext variants in the Pareto front, both SHA-3 and ad hoc checksums show quite often. CRC8 shows only seldomly. CRC8 and BCH are often doubles, i.e., they have the same distance and bandwidth as DYST-Basic or -Ext with SHA-3. Quite some variants from the theoretical analysis are not on the Pareto front, indicating that sometimes the simulations gain a little in practice.

D. Throughput-Optimized DYST Using Multiple Pointers

For DYST-Basic and DYST-Ext, only one hash is calculated per PoI and eventually pointed to. DYST’s throughput can be increased by using multiple pointers, signalling that a secret message might be found after calculating one of multiple hashes using a counter i , so that $H_1 = H(M||i = 0), \dots, H_n = H(M||i = n - 1)$, i.e., the signal (if sent) tells the receiver what counter i needs to be used to obtain the secret message. This requires n pointers in the form of n distinguishable inter-packet times between PoI and signaling packet (or a classical covert storage channel with $\log_2 n$ pointer bits).




Fig. 7. History and prediction covert channels, differentiated by the secret data carrying transmission they *point* to.



		Covert Sender		
		Active (generates own overt traffic in which it embeds covert data)	Passive (embeds covert data in overt traffic of third-party nodes)	Fully-passive (utilizes third-party traffic without modifying it)
Covert Receiver	Active (is the destination of the overt traffic)	Active Covert Channel	Semi-passive Covert Channel	Fully-and-semi-passive Covert Channel
	Passive (is not the direct destination of the overt traffic, e.g., a router)	Semi-active Covert Channel  DYST's Signal Channel	Passive Covert Channel	Fully-passive Covert Channel  DYST's Data Channel

Fig. 8. Categorization of DYST's data and signal channels.

E. Taxonomy

Existing publications on network covert channels exclusively focus on network traffic that is live traffic to be modified or generated. Some covert channels also replay traffic recordings enhanced with secret data. *History* covert channels *point* to secret data in a carrier that was transmitted in the past (Fig. 7), i.e., the carrier traffic of the data signal is not altered. While not a core aspect of this paper, we still like to point out that it would be imaginable to create *prediction* covert channels, which point to anticipated future data. For instance, ARP requests in LANs and sensor value readings in CPS occur on a regular basis and are thus possible to predict and utilize. The only difference between prediction and history covert channels is whether they point to old or upcoming data. However, predictions of future traffic are less reliable than pointing to already-seen traffic. For this reason, we solely focus on history channels.

The literature differentiates covert channels into active and passive ones (Fig. 8). An active sender generates the traffic in which the secret data is embedded while a passive sender modifies third-party traffic for this purpose. Usually, a passive sender is an intermediate network node, such as a router. The receiving process can also be performed in an active or passive manner. Here, the terminology considers a receiver as active if it is also the destination of the overt traffic. If it passively observes the traffic (which is directed to another hop), the receiver is considered passive.

Lamshöft and Dittmann recently added a further differentiation in [57], which is also shown in Fig. 8. They consider covert channels as semi-active if the covert sender is active but the covert receiver is passive. In contrast, a channel is called semi-passive if the covert sender is passive but the covert receiver active.

As the current differentiation between active and passive covert channels does only represent the *signaling* channel of

DYST, we add a new category of covert communication, which we call *fully-passive* because of its truly passive handling of third-party traffic (which is not modified). Because of the broadcast nature of the utilized messages, the receiver of DYST's *data* channel is a passive one.

As our channel's sending and receiving processes are decoupled in a way that the sender does not directly address the receiver, our data channel can further be considered as an *indirect* one, while the signaling channel can be considered as a *direct* covert channel.

Finally, a covert channel could also be a *fully-and-semi-passive* one, which is at least theoretically feasible and reflects a channel where the covert receiver waits for pre-defined packets directed to it by some third-party node (fully-passive covert sender). Such a channel could be configured by having DYST operate with directed messages instead of broadcast messages and would be less functional.

IV. EVALUATION USING LOCAL NETWORK SETUP

This section presents the evaluation of the DYST implementations Basic and Ext for local area networks using two different scenarios: a private smart home network and a university network. After presenting our implementation and the experimental setups for different scenarios in which we evaluated DYST, we analyze the robustness and detectability of our covert channel.

A. Implementation

The PoC for DYST using local area networks (DYST-Basic/Ext) was implemented with Python 3 and utilizes the *scapy* library for eavesdropping legitimate traffic and crafting signal packets. Our implementation utilizes the following packets as they can be received by CS and CR when residing in the same network:

- IPv6 anycast packets with IPv6 destination `ff0* : :` or to IPv6 link layer address `33 : 33 : *`
- IPv4 broadcasts to the subnet broadcast address
- ARP requests to broadcast addr. `ff : ff : ff : ff : ff : ff`

For hashing, we utilized the SHA3 hash algorithm with a bit length of 512, provided by the Python 3 library *hashlib*. The input values contained the source IPv4 and IPv6 address, depending on the packet type. As the same input of a hash function results in the same hash, we additionally utilized the CS and CR packet receiving timestamp in seconds. The utilization of the timestamp results in a new hash for the same source addresses each second. As packets are not received at the same time by all devices in a network, we filtered packets that were received closer than 0.05 seconds to a full second, i.e., having a fractional time value of less than 0.05 or higher than 0.95 seconds, to minimize the possibility of CS and CR using a different timestamp when hashing.

For signaling, we utilized an ARP broadcast request, sent by the CS requesting the MAC address containing the target IPv4 address of an uninvolved third-party system. The CR interpreted this request as the signal to extract the latest hashed value. The process of the CS implementation is shown in Algorithm 1. An additional check for PoI-collisions is performed with and without activated robustness mode (cf. Section IV-D).

Algorithm 1: DYST CS-Side Implementation.

```

Input: Secret message  $M$ 
Split secret message  $M$  into chunks  $chunk_0, \dots, chunk_n$ 
 $desired\_hash = chunk_0$ 
for Observed packet in traffic do
  if packet is PoI then
     $hash = H(packet||timestamp)$ 
    if  $hash == desired\_hash$  then
      if NOT Robustness Mode then
        Send signal
      end
    else
      if No collision within timespan  $D$  then
        Send signal
         $desired\_hash = chunk_{next}$ 
      end
    else
      Ignore packet
    end
  end
end
else
  Skip packet
end
end
else
  Ignore packet
end
end

```

B. Scenarios and Testbed

To evaluate DYST under different circumstances, we came up with several scenarios which provide different traffic characteristics. These scenarios are described in detail in the following paragraphs.

a) Scenario 1. University Network: Traffic for this scenario was recorded in a university network from regular office workstations. We did not use any port mirroring or a prominent location in the network to see how a regular device would see traffic. The environment itself is composed of 75 to 100 devices, around 50 of which are used on a daily basis. The network mostly consists of office laptops, printers, and some smart devices. All major operating systems (Linux, macOS, Windows) are present. For our intents and purposes, the university network resembles that of a company, thus this scenario applies to use cases in both settings.

Example use case: In an APT, an attacker might infect multiple clients in the network and use DYST as a means of internal communication between the compromised clients. If one infected machine got access to an account with higher privileges, it could share the credentials with all other instances in the network for a faster spread of malware. Similarly, DYST could be used as a command and control channel between multiple compromised clients.

b) Scenario 2: Home Network: This scenario represents a typical home network with mixed devices, permanently connected to a WiFi router. The utilized router was a Speedport Smart 3 with current firmware, extended by two mesh repeaters to cover a larger area. In total, up to 30 devices were connected simultaneously, consisting of classical IT devices (three laptops, two raspberries, a network printer, smartphones), IoT devices




Fig. 9. Match distribution.

(SmartTV, vacuum cleaning robot, coffee machine) as well as home automation (various Google Nest Mini). All devices were commonly used and, except for laptops and smartphones, connected permanently to the home network. All devices were connected within one /24 IPv4 subnet.

Example use case: Several compromised smart home devices exchange information under the radar, e.g., to collaboratively collect surveillance data and profile inhabitants. It is not probable that such a network is monitored for covert channel detection, so a less sophisticated approach will also be applicable. Anyhow, we decided to analyze this scenario because it shows the flexibility of DYST, even if there are few changing devices. In this scenario, especially the throughput of DYST can be optimized as there are no wardens.

Note that further scenarios are imaginable, e.g., journalists exchanging secret information through a WiFi hotspot in a bullet train or at a public airport.

C. Match Distribution

The utilized input parameters of the hash function generate different hashes h_i for each modified bit, i.e., new packet, if the hash function is collision-resistant. These generated hashes h_i are compared by the CS to a specific pattern M (the data it wants to signal), which is constant until a suitable match is found. As explained in the derivation of (5), the number of matching bits follows a binomial distribution with h trials and a success probability 0.5.

We compare the actual frequency for the number of matching bits in relation to the total number of hashes generated to the expected results in Fig. 9 for both scenarios. Therefore, we searched for one specific pattern equal to 8 bits and expected the distribution to follow the binomial distribution with $N = 8$ and $p = 0.5$, which is represented by the gray histogram. Further, we assume that the more hashes are observed, the more the actual distribution will follow the expected distribution. The black and red bars represent the actual distribution of matching bits in the home and university network scenarios, respectively. Both actual distributions follow the expected distribution, however, the university network scenario differs slightly. This can be explained by the number of observed hashes (89,959 observed hashes for the home network; 2,075 for the university network), drawing both assumptions correct. This points out our hash-generation methodology is correct.

TABLE III
EXPECTED AND ACTUAL MATCH RESULTS (HOME NETWORK)

h	8 bits	12 bits	16 bits	21 bits
# observed packets	89,959	56,105	226,488	43,325
$P_h(X \geq h)$	0.3125%	0.2930%	0.0024%	0.0001%
freq. pkts w. 100% match	0.3813%	0.2673%	0.0022%	0.0000%
$P_h(X \geq 0.8 \cdot h)$	3.5937%	1.9360%	1.0504%	0.3917%
freq. pkts w. $\geq 80\%$ match	3.5160%	1.9035%	1.0654%	0.3508%

Besides the evaluation in our testbeds, also the number of matches for different bit lengths is evaluated. As the CS is searching for a 100 percent match in DYST-Basic, we calculated $P_h(X \geq h)$. As for DYST-Ext also partial matches can be utilized, we add $P_h(X \geq 0.8 \cdot h)$ for at least 80% matching bits, both calculated by (5).

The results of the example experiments for the home network scenario are presented in Table III. The 8-bit and 16-bit experiments were performed in a live scenario, while the 12-bit and 21-bit match distributions were simulated. The results indicate that for $h = 8$ bits, slightly more hashes than expected show h matching bits, while for at least 80% matching bits, slightly less were observed. According to the deviation, 281 perfect matches should be detected, while actually 343 were detected. For the DYST-Ext mode, 3,232 potential matching hashes should be detected, while actually 3,162 were detected. For our $h = 12$ bit experiment, the hit rate for both 100 and (at least) 80 percent matches slightly performed worse than expected, resulting in 150 matches instead of 164 and 1,068 instead of 1,086, respectively. For $h = 16$ bit and a 100 percent match, the observed rate neared the expected value, while the number of actual 80 percent matches was slightly higher than expected. There should have been 5 matches, while actually 5 had been found for DYST-Basic and for DYST-Ext potential 2,379 hashes with at least 80 percent match should be observed, while 2,413 were actually detected. For our $h = 21$ bit experiment, no matches were found. As 0.43 packets are expected for the number of observed packets, the expectation is met. Further, 152 packets had at least 80% matching bits, while there should have been 170 packets.

D. Robustness

The main concern for the robustness of DYST lies on the question which messages are seen and interpreted for signaling. As DYST only uses legitimate network packets, it relies on the general robustness of network transmissions (e.g., Ethernet frame checksums) and timestamps (ensured by time synchronization mechanisms like NTP for example). Such effects are caused by active and inactive jitter-influencing factors causing disorder, retransmissions, dropouts, and delays amongst others. Neither CS nor CR can control jitter factors of their transmission as they are not *directly* communicating with each other. To mitigate such factors caused by the physical layer we need to introduce the so-called *robust mode*.

We evaluated the jitter in our two local network environments for DYST-Basic by investigating the delays between ARP requests and the corresponding replies. Fig. 10 shows the results for the university and home network.




Fig. 10. ARP delays for University and Home network setups.

We can see that the mean values and standard deviations for the ARP delays, which are by nature two-way-delays, are orders of magnitude lower than the waiting periods that were used for the robustness measure (see further below). The standard deviation for the university is at $9.1 \cdot 10^{-3}$ s and for the home network at $2.1 \cdot 10^{-5}$ s. While the university network has generally higher values, we still only have 35 outliers above 0.1 s out of over 18,500 measurement points. Therefore, we can conclude that the jitter of our test networks did not have a relevant impact on the robustness of DYST. Further, it is also viable to adjust the waiting periods for sending signal packets according to the characteristics of the network in which one chooses to deploy DYST.

If the same packet arrives at CS and CR, we can assume that the content will be the same, resulting in the same input parameters for the hashing function. We do have other concerns about the robustness of DYST that need to be addressed:

- a) CS and CR must receive the same PoI.
- b) CS and CR must receive the PoIs in the same order.
- c) Consecutive PoIs must have a sufficient delay to allow reliable signaling.

a: To address this issue, we have to carefully select which packets are used by DYST to ensure that both CS and CR receive the same packets. Depending on the deployment scenario, we might choose different sets of packets to achieve this goal: In a local scenario like a university network, a smart home network or an open Wi-Fi like a café, we can focus on local broadcast packets. If DYST is used between two routers, we can use our knowledge about the routing topology to filter for packets that will pass through both CS and CR. Similarly, if both CS and CR are part of the same multicast group, one could filter for packets from that group to ensure synchronization. Moreover, there is a tolerance regarding the received PoI. As only a fraction of the packets are actually used for signaling, a packet received by CR but not received by CS would simply not be checked for a match and thus also not used as a carrier. However, if CS receives a packet and both of the following conditions are met: CR does not receive the packet and CS actually points to this packet for signaling, then CR might receive an incorrect message as it interprets the wrong packet (see point b)). All in all, it is possible to choose a robust set of packets to be used with DYST, with only a little prior knowledge about the deployment scenario.

b: This variable is outside the control of DYST as we cannot influence the routing or buffering behavior of other parts of the




Fig. 11. Robustness measures.

network. We performed an evaluation in a home network by running DYST between two different clients on the same network. During our evaluation, we observed significant problems when testing our scripts. In our first test, 3 of 65 characters of the message were transmitted correctly. This was due to the fact that a significant portion of the PoIs did not arrive in the same order for the CS and the CR. These PoI packets were in the wrong order because of their tight succeeding timing and different networking delays for CS and CR. Therefore, two additional configuration parameters for DYST control the mandatory delay between received packets to reduce this issue: If the CS receives two or more PoIs in less than D milliseconds, the CS will ignore all PoIs received in that timeframe. So only isolated PoIs will be considered for DYST. If only one PoI is received in D milliseconds, the CS will send out the signal. The CR will ignore all PoIs that arrived less than R milliseconds before the signal and only interpret earlier PoIs. This gives the CS enough time to calculate and send the signal without risking a race condition (see part c) below). Fig. 11 illustrates this process.

Fig. 11(a), shows the side of the covert sender. This example uses two packets that arrive close to each other, which leads to them being ignored by the CS and not considered for DYST. Fig. 11(b) shows the receiver-side. Here we can see that the receiver ignored a packet that arrived too close to the signal and instead interpreted the older one. This will decrease the potential throughput, as we ignore more packets that could potentially be useful. But we significantly decrease potential errors in cases where packets arrive in a different order at the CS as at the CR. Varying scenarios will require differently tuned values for D and R , see Electronic Supplement. Additionally, we can counter possible errors with an error-correcting code that covers multiple transmissions. Different approaches are possible, e.g., simply transmitting the same message multiple times and taking a majority vote on gained hash values, can increase robustness.

c: Similarly to b), DYST might encounter errors if many PoIs arrive in a short amount of time. The CS might receive a PoI that creates a hit. While the CS is evaluating and preparing to send out the signal, another PoI arrives at the CR even before the signal from the CS reached the CR. In such a case, the CR would interpret the latest PoI and not the correct one. This error source is also countered by the option R of the multistage delay introduced in b). Since DYST only considers PoIs that are isolated and the CS has R milliseconds to perform the signaling, it is far less likely for this race condition to appear for any PoIs that can be seen by CS and CR. Similar to b), a larger delay will result in a lower bandwidth but higher robustness. Depending




Fig. 12. Comparison of inter-packet delays (IPDs) between PoIs (SIMULATED).

TABLE IV
ROBUSTNESS EVALUATION - MATCHING TRANSMITTED CHARACTERS

	Non-robust Matches	Robust Matches
University Network	100%	100%
Home Network	4%	56%

on the scenario, drastically different configurations for D and R are possible. If CS and CR are both routers with a fixed route between them, it is easy to see that the CS will have better knowledge about the order and delays in which the CR will receive the PoIs and can therefore choose a lower value for D and R .

Impact of Robustness: We evaluated the timing between PoIs in both scenarios of our local network testbeds to get an overview of network behavior and the impact of our robustness measure. Fig. 12 shows the plots for the different scenarios.

We can see that both scenarios show a generally similar picture: the mean value for the inter-packet delays (IPDs) of the PoIs for DYST-Basic is close to 0 with a significant number of outliers. The outliers are beneficial for the robustness of DYST, as they separate PoIs. If we look at the delays with the robustness filter ($D = 0.5$ s, the simulation only happened for the CS side), we again see a similar image for all scenarios. We generally see higher delays between PoIs, as DYST ignores some PoIs and therefore the delays between evaluated PoIs are higher. This means that DYST will have fewer packets at its disposal to transmit the message, but we gain in reliability (see Electronic Supplement for details).

We evaluated the effectiveness of our robustness approach in the home network and university network scenarios. Table IV shows the results for two scenarios (home and university network). For each scenario, we conducted a test run with and without robustness measures and recorded the percentage of characters that were correctly transmitted.

We can see that the home network had significant problems without robustness measures, while the university network setup had no problems during our tests. The university network setup showed very little activity, i.e., very low frequency of PoI, during our robustness tests, which resulted in a reliable transmission but also in a low absolute bandwidth (8 characters in 8 hours) despite usual frequency of matching PoI. The home network setup had drastically more activity and higher bandwidth (48 characters in 7 hours), which in turn resulted in a significantly less reliable transmission. Initially, we tested with $D = 0.3$ s

TABLE V
ROBUSTNESS EVALUATION - TOTAL POIS (SIMULATED)

	Non-robust	Robust ($D = 0.5s$)	Fraction (%)
University Network	8,130	4,595	56.5%
Home Network	79,319	12,689	15.9%

TABLE VI
OVERALL RECORDED TRAFFIC

		# Pkts	# ARP Requests	Rec. Time
Legitimate	Home Network	17,608,213	8,619,952	40 Days
	University Network	2,096,387	218,707	15 Days
Covert	Home Network	5,604,205	3,023,422	11 Days
	University Network	5,380,286	522,451	56 Days

and $R = 0.015$ s which resulted in ca. 31% correctly transmitted characters. We further tested the parameters $D = 0.5$ s and $R = 0.3$ s, which resulted in 56% correctly signaled characters. We can therefore see that the home network setup benefits from more aggressive configurations. However, the robustness would increase using a higher D or additional robustness measures, such as the redundant transfer of secret messages.

In Table V, we show how many PoIs were observed in a simulated offline run. We used real recordings and ran the pcaps through an offline version of DYST (only CS side is simulated). We can see that in all scenarios, the number of usable PoIs is significantly reduced. This means there will be fewer PoIs available for DYST, which in turn can reduce the potential bandwidth of the covert channel. It is noticeable that the home network setup suffered more than the university network setup, this can be explained by the higher activity in the home network compared to the university network. This again explains the worse performance in the home network without robustness measures.

On the other hand, a reduction of possible bandwidth aids the undetectability, by spreading signals even further apart.

E. Detectability

Detectability of a timing channel is analyzed by comparing the distributions of inter-packet gaps of packet streams with and without timing channel. This can be done with statistical and information-theoretic means. We have used Kolmogorov-Smirnov-test for the former and Cabuk's compressibility score for the latter. The results of these tests are not absolute: the timing channel can only be detected if the difference between results with and without timing channel is *larger* than the variation of test results for different packet streams without timing channel.

To evaluate the detectability of DYST, we gathered legitimate reference and covert channel recordings from two different scenarios and with six different configurations for the covert channel: 1 and 2 B basic mode (h has 8 or 16 bits, respectively), 1 and 2 B robust mode (also 8 and 16 bits), and 1 and 2 B extended mode (where h is also 8 or 16 bits, respectively. The checksum of length c is added on top). We performed recordings in two different networks, a home network and a university network (see Table VI for an overview).

As the data channel itself cannot be detected, a defender relies on detecting the signal channel. Since our DYST implementation uses ARP requests for signaling, our detection focuses on the IPDs of ARP requests.

1) *KS-Test*: To gauge the potential detectability, we chose a two-sample KS-test [58]. The KS-test is a general measure of similarity between two samples, with high test results indicating difference of the samples [59]. Given empirical cumulative distribution functions F_1 and F_2 for random variables X_1 and X_2 , i.e., $F_i(a) = P(X_i \leq a)$, the similarity of the two functions is computed by $\sup_a |F_1(a) - F_2(a)|$, with \sup being the supremum of the set of distances over all real values a [59].

We use the KS-test to measure the similarity between the ARP IPD distribution of two recordings. In addition to our original recordings, we also filtered the covert channel recordings to remove any signals produced by the covert channel to produce a second (but synthetic) source of legitimate traffic. We then performed a cross validation between all possible recordings of a scenario (home and university networks). This provides us with 3 different classes of combinations, which we considered important for our analysis: 1) covert versus filtered recordings, 2) pairs of legitimate recordings, and 3) covert versus legitimate recordings. To get a better understanding of the detectability, we focused on several combinations, which we describe separately.

First, we compared different *legitimate* recordings against each other, while excluding exact matches. For this, we used several legitimate recordings from the home network. This gives us a mean p-value of $3.84 \cdot 10^{-11}$ with a standard deviation of $4.01 \cdot 10^{-10}$ and a mean D-value of 0.19 with a standard deviation of 0.13. The results for the legitimate university network scenario are almost exactly the same. This points towards a significant difference between all the legitimate recordings. With that, we can already see that legitimate traffic drastically varies depending on the time of day and the activity of participating nodes that do not follow repetitive behavior. This already points to a low possibility of detection.

Next, we compared the *covert channel* recordings with the corresponding *filtered covert channel* recordings. This gives us a mean p-value of 0.99 with a standard deviation of 0.04 and a mean D-value of 0.0008 with a standard deviation of 0.001. This on the other hand shows that our covert channel barely alters the characteristics compared to legitimate traffic.

If we compare *covert channel* recordings with *legitimate* recordings, we obtain a mean p-value of $1.53 \cdot 10^{-95} \sim 0$ with a standard deviation of $1.12 \cdot 10^{-94} \sim 0$ and a mean D-value of $1.87 \cdot 10^{-1}$ with a standard deviation of $1.31 \cdot 10^{-1}$. These results again point towards significant differences between the two scenarios. However, these differences are comparable to those of two legitimate recordings.

Fig. 13(a) shows the p-values of the KS-tests and we can see a strong similarity of the CC versus filter scenario and a strong dissimilarity for the other two scenarios. Fig. 13(b), (c) and (d) show exemplary histogram plots for the ARP request IPDs for each of the considered scenarios from the home network setup. Again, we can determine a high level of similarity between the covert channel recording and the filtered recording and slight differences between two legitimate recordings as well as

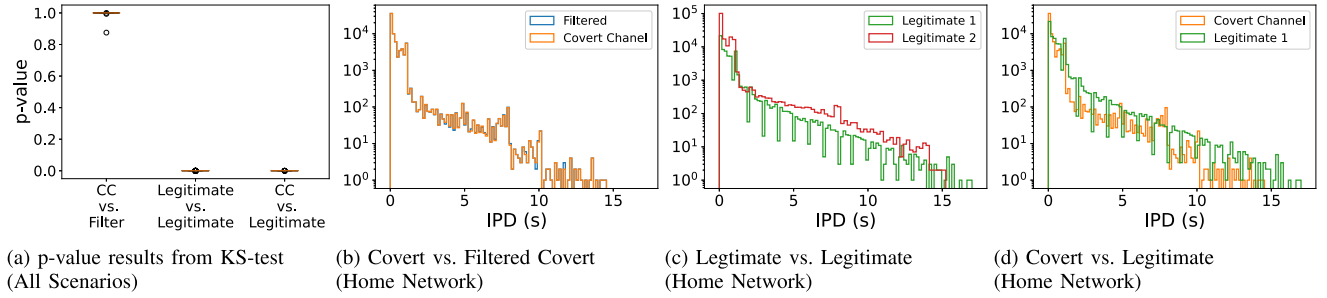


Fig. 13. Overview of legitimate and DYST traffic's characteristics, with detailed examples for IPD values of the home network.

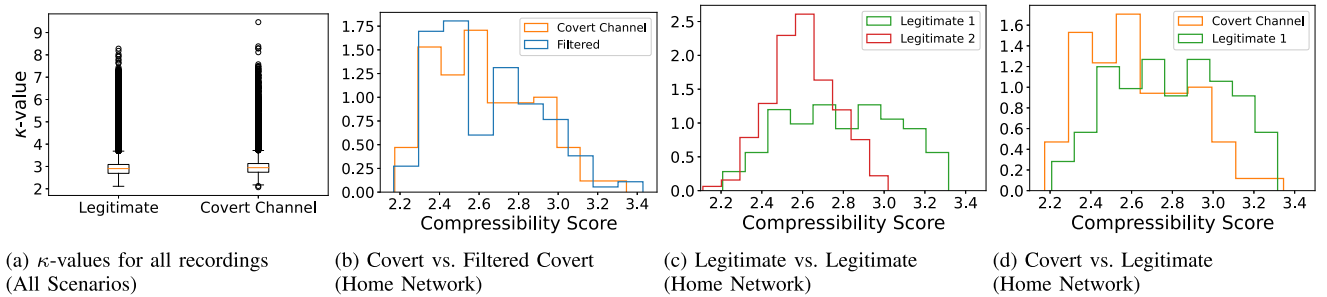


Fig. 14. Overview of legitimate and DYST traffic's compressibility scores, with detailed examples for the home network scenario.

between covert and legitimate recordings. These three results in combination point towards an almost impossible detection of DYST, as the differences between different legitimate recordings are larger than the difference between covert channel and filtered recordings.

2) *Compressibility Score*: In addition to the KS-test, we also used a widely known detection method from covert channel research: the compressibility score as proposed by Cabuk et al. [25]. Again, we used the IPDs of ARP requests for the detection and focused on the same three classes of combinations with the same recordings. To calculate the compressibility score, we divided the recordings in windows of fixed length (1,000 IPDs). Each window was then transformed into a string representation S of concatenated IPDs, which was compressed using a compressor \mathfrak{S} , in our case `gzip`. The final compressibility score $\kappa = |S|/|\mathfrak{S}(S)|$ for a window is the compression ratio between the original string and the compressed string. Since covert timing channel flows use similar IPD values in a re-occurring manner, their compressibility score is typically higher than the score of legitimate flows.

Fig. 14(a) shows the compressibility scores for all recordings. Here we can see a nearly identical values for legitimate and covert channel recordings, which already points towards a challenging detectability of DYST. Fig. 14(b), (c) and (d) show exemplary histograms with pairwise plots for the distributions of the compressibility scores of the home network. As can be seen, we found a large overlap of compressibility scores when comparing covert channel and filtered recordings. This again points towards a strong similarity of the two recordings. The comparison of two legitimate recordings of the home network shows significant differences in the distributions while we would expect two

matching plots, indicating a high dependence on the current status of a network (daytime etc.). Moreover, the comparison between the covert channel and the legitimate recording shows even smaller differences than the two legitimate recordings. When determining optimal thresholds for the detection of DYST, our experiments revealed that the compressibility metric (and the related AUC scores for ROC charts) fluctuated based on time of day, network load or other factors, rather than the presence of DYST. This is rooted in the fact that DYST sends only very few messages. Thus, similar to the KS-test, we observed that there is no clear threshold for the κ -value to discriminate between legitimate and covert channel recordings. This fact leads to the conclusion that DYST is not detectable with the compressibility score.

3) *Evaluation of Different Covert Channel Configurations*: We also evaluated the detectability of DYST when only focusing on a single configuration at a time (e.g., 1 B basic or 2 B extended modes). We found that the compressibility score performed no different than before, still producing overlapping κ -values. Similarly, the KS-test performed comparable (figures omitted due to space reasons). This was expected, as the detectability of DYST relies primarily on the number of signals sent.

To further evaluate detectability, we ran our multi-pointer variant of Section III-D while allowing different pointer counts. Fig. 15 shows the detection results using the compressibility score.

As can be seen, the approaches with two, and even 32 pointers perform very close to DYST-Basic with only one pointer (AUC close to 0.5). Only when the number of pointers was increased to ≥ 128 , we were able to provide reasonably good detectability that would be useful in practice.




Fig. 15. Detection performance based on pointer count.

TABLE VII
AVG. REMOTE TRANSMISSION TIMES (13–15 CHARACTERS)

Scenario	Robust	From First Signal	From Start
Home Network (local)	False	00:26:55	00:30:53
	True	03:15:08	03:38:49
University (local)	False	01:11:27	01:28:14
	True	03:22:14	04:06:29
Remote-Smarthome	False	00:10:13	00:11:40
	True	08:06:28	09:46:30

We conclude that DYST’s detection depends on the chosen configuration. A configuration with ≤ 32 pointers (i.e., high CAF, but few matches and low bitrate) was undetectable. A configuration with ≥ 64 pointers sacrifices is detectable and offers a low CAF but provides a higher bitrate (cf. Section VI).

V. FEASIBILITY STUDY OF REMOTE VARIANT

So far, we analyzed DYST-Basic and DYST-Ext in a local network. We now study the feasibility of DYST-Remote-Smarthome. In particular, we use a traffic recording from a smart home testbed to evaluate the sending performance of DYST with three different short messages and the number of PoIs that were observed during these tests.

Performance Comparison. Local versus Remote: We utilized traffic recordings from our university’s home automation testbed running eleven different smart devices, which are connected to a local router with NAT-based Internet access. The recordings were performed over a duration of several weeks while the network was in *idle* mode, i.e., no user was interacting with these devices. To simulate the sending process of DYST-Remote-Smarthome, we read the pcap-recordings packet by packet and performed the same calculations as would be done by a real CS. When a match was found, instead of sending out a signal packet, we wrote a corresponding line to a logfile.

To compare the performance of DYST-Remote-Smarthome with DYST-Basic, we transmitted three short messages (13–15 characters) using different configurations. Table VII shows the average time it took to transmit a message (transmitting 1 B blocks, i.e., $h = 8$) for a given scenario, with and without robust mode.




Fig. 16. Bandwidth evaluation results for the multi-pointer approach and two different recordings from the office network (different times of the day).

As we can see, DYST-Remote-Smarthome delivers better transmission speeds without the robust mode while the transmissions with the robust mode are noticeably slower compared to the other two scenarios. This is to be expected, as we found more potential PoIs in the remote scenario (12,700 versus 7,300 and 2,100 PoI per hour), compared to the local mode which results in more potential matches (non-robust mode), but also more potential collisions (robust mode).

Further Notes: The robustness measures and the optimization measures for DYST-Remote-Smarthome work the same way as for the LAN scenario; the major difference relies solely on the type of packets that are considered PoI and the type of packets that are signaling packets. For this reason, and because of the simulation-only data, we skip an evaluation of robustness and optimization at this point. The detectability of DYST-Remote-Smarthome can be considered similarly infeasible as with DYST-Basic as we only sent 1,345 signaling packets per day (on average, $\sigma = 364$), and even fewer messages while in robust mode.

VI. THROUGHPUT OPTIMIZATION THROUGH MULTI-POINTER METHOD

The previous experiments focused on a DYST setup with maximized amplification (CAF). We additionally evaluated the throughput-optimized setup that was introduced in Section III-D. To evaluate this approach against the performance of a traditional covert storage channel, we implemented a synthetic version that works on PCAP recordings and analysed the behaviour on two recordings from a busy office network. The results are shown in Fig. 16(a) and (b). The *direct embedding* represents a covert storage channel that instead of embedding the pointer embeds the data itself. Note that we chose to only embed covert data in packets that would have been used as signal packet by DYST. As can be seen, we can increase the sending performance by sacrificing the CAF up to a point where we have no amplification but the full bitrate of a normal covert storage channel. In other words, one can configure DYST so that one either optimizes for stealthiness (maximize CAF) or throughput. Note that the CAF becomes zero if the pointer size equals the message size. We can observe that for both scenarios (i.e., pointing to either 1 or 2 bytes per match), DYST has the advantage over the direct embedding up to the point where the

pointer size matches (or exceeds) the size of the message being pointed to. The maximum bitrate *with* amplification was 29.12 (7 b pointer to 8 b message) and 58.75 (15 b pointer to 16 b message), respectively.

VII. DISCUSSION

Our approach provides room for several follow-up developments in terms of alternative methodology. We have studied how one can point to the *content* of previously seen packets. However, other hiding patterns [10] could probably be useful, too. For instance, the *inter-packet times* hiding pattern could be used to point to a series of inter-packet gaps, the *size modulation* pattern could be used to point to a sequence of packet sizes and so forth. Further, by neglecting the context of network traffic and discarding the use of a hash function, we have recently shown that a history channel can be used to point to textual elements on the Internet (e.g., content from Wikipedia pages) to bypass censorship [60]. A similar scenario could be a social media service, where legitimate users submit postings, such as *Facebook*, *Twitter*, or any form of blog or platform that generates a massive amount of content (e.g., *Dropbox*, *Google Drive* or *Github*). Signaling for these scenarios is not bound to the same service, which means that signals can be sent out-of-band, for example via a different platform or a network protocol. Moreover, one could use history covert channels to point to previously seen content in audio/video streams. Finally, DYST could be applied for stealthy communication between local processes of a secure operating system by monitoring hardware events.

VIII. CONCLUSION

We presented the history covert channel method. History covert channels send secret messages by *pointing* to unaltered legitimate data. They can be applied in extremely hostile environments, i.e., where a covert sender is unable to send any other message than previously whitelisted ones. We presented two feasible encoding strategies and implemented DYST in different network environments. Our results indicate a limited but variable throughput dependent on the number of bits signaled at once. We analyzed the robustness and also have shown that traditional heuristics are unable to provide satisfying results unless many pointers are employed. Future work will focus on alternative on improving the sending performance and developing additional countermeasures. Code and data are available: <https://github.com/NIoSaT/DYST>

REFERENCES

- [1] B. W. Lampson, "A note on the confinement problem," *Commun. ACM*, vol. 16, no. 10, pp. 613–615, 1973.
- [2] DoD, "Trusted computer system evaluation criteria," in *The 'Orange Book' Series*, London, U.K.: Palgrave Macmillan, 1985, pp. 1–129.
- [3] F. A. P. Petitcolas, R. J. Anderson, and M. G. Kuhn, "Information hiding-a survey," in *Proc. IEEE*, vol. 87, no. 7, pp. 1062–1078, Jul. 1999.
- [4] S. Zander, G. J. Armitage, and P. Branch, "A survey of covert channels and countermeasures in computer network protocols," *IEEE Commun. Surveys Tuts.*, vol. 9, no. 3, pp. 44–57, Third Quarter, 2007.
- [5] B. Carrara and C. Adams, "Out-of-band covert channels - A survey," *ACM Comput. Surv.*, vol. 49, no. 2, pp. 23:1–23:36, 2016.
- [6] L. Cavaglione and W. Mazurczyk, "Never mind the malware, here's the stego-malware," *IEEE Secur. Priv.*, vol. 20, no. 5, pp. 101–106, Sep/Oct. 2022.
- [7] L. Cavaglione, "steg-in-the-wild," 2023. [Online]. Available: <https://github.com/lucacav/steg-in-the-wild>
- [8] D. Barradas and N. Santos, "Towards a scalable censorship-resistant overlay network based on WebRTC covert channels," in *Proc. Int. Workshop Distrib. Infrastructure Common Good*, 2020, pp. 37–42.
- [9] A. Mileva and B. Panajotov, "Covert channels in TCP/IP protocol stack - Extended version-," *Central Eur. J. Comput. Sci.*, vol. 4, no. 2, pp. 45–66, 2014.
- [10] S. Wendzel, S. Zander, B. Fechner, and C. Herdin, "Pattern-based survey and categorization of network covert channel techniques," *ACM Comput. Surv.*, vol. 47, no. 3, pp. 50:1–50:26, 2015.
- [11] S. Wendzel, W. Mazurczyk, and G. Haas, "Don't you touch my nuts: Information hiding in cyber physical systems," in *Proc. IEEE Secur. Privacy Workshops*, 2017, pp. 29–34.
- [12] P. Krishnamurthy, F. Khorrami, R. Karri, D. Paul-Pena, and H. Salehghafari, "Process-aware covert channels using physical instrumentation in cyber-physical systems," *IEEE Trans. Inf. Forensics Secur.*, vol. 13, no. 11, pp. 2761–2771, Nov. 2018.
- [13] M. Hildebrandt et al., "Threat analysis of steganographic and covert communication in nuclear I&C systems," in *Proc. Int. Conf. Nucl. Secur.: Sust. Strength. Efforts*, 2020, p. 14.
- [14] K. Lamshöft, T. Neubert, C. Krätzer, C. Vielhauer, and J. Dittmann, "Information hiding in cyber physical systems: Challenges for embedding, retrieval and detection using sensor data of the SWAT dataset," in *Proc. ACM Workshop Inf. Hiding Multimedia Secur.*, 2021, pp. 113–124.
- [15] J. K. Millen, "20 years of covert channel modeling and analysis," in *Proc. IEEE Symp. Secur. Privacy*, 1999, pp. 113–114.
- [16] Q. Xu, H. Naghibijouybari, S. Wang, N. B. Abu-Ghazaleh, and M. Annavaram, "GPUGuard: Mitigating contention based side and covert channel attacks on GPUs," in *Proc. Int. Conf. Supercomput.*, 2019, pp. 497–509.
- [17] M. Urbanski, W. Mazurczyk, J. Lalande, and L. Cavaglione, "Detecting local covert channels using process activity correlation on Android smartphones," *Int. J. Comput. Syst. Sci. Eng.*, vol. 32, no. 2, pp. 71–80, 2017.
- [18] K. Block, S. Narain, and G. Noubir, "An autonomic and permissionless android covert channel," in *Proc. Conf. Secur. Privacy Wireless Mobile Netw.*, 2017, pp. 184–194.
- [19] M. Hanspach and M. Goetz, "On covert acoustical mesh networks in air," *J. Commun.*, vol. 8, no. 11, pp. 758–767, 2013.
- [20] M. Guri, M. Monitz, Y. Mirski, and Y. Elovici, "BitWhisper: Covert signaling channel between air-gapped computers using thermal manipulations," in *Proc. IEEE 28th Comput. Secur. Found. Symp.*, 2015, pp. 276–289.
- [21] M. Guri, B. Zadov, and Y. Elovici, "LED-it-GO: Leaking (A lot of) data from air-gapped computers via the (small) hard drive LED," in *Proc. Int. Conf. Detection Intrusions Malware Vulnerability Assessment*, Springer, 2017, pp. 161–184.
- [22] M. Guri, Y. A. Solewicz, and Y. Elovici, "MOSQUITO: Covert ultrasonic transmissions between two air-gapped computers using speaker-to-speaker communication," in *Proc. IEEE Conf. Dependable Secure Comput.*, 2018, pp. 1–8.
- [23] C. Gong, X. Yue, Z. Zhang, X. Wang, and X. Dai, "Enhancing physical layer security with artificial noise in large-scale NOMA networks," *IEEE Trans. Veh. Technol.*, vol. 70, no. 3, pp. 2349–2361, Mar. 2021.
- [24] L. Mucchi, S. Caputo, P. Marocci, G. Chisci, L. Ronga, and E. Panayirci, "Security and reliability performance of noise-loop modulation: Theoretical analysis and experimentation," *IEEE Trans. Veh. Technol.*, vol. 71, no. 6, pp. 6335–6350, Jun. 2022.
- [25] S. Cabuk, "Network covert channels: Design, analysis, detection, and elimination," Ph.D. dissertation, Purdue Univ., West Lafayette, IN, 2006.
- [26] S. Cabuk, C. E. Brodley, and C. Shields, "IP covert channel detection," *ACM Trans. Inf. Syst. Secur.*, vol. 12, no. 4, pp. 22:1–22:29, 2009.
- [27] S. Cabuk, C. E. Brodley, and C. Shields, "IP covert timing channels: Design and detection," in *Proc. 11th ACM Conf. Comput. Commun. Secur.*, 2004, pp. 178–187.
- [28] G. Shah and A. Molina, "Keyboards and covert channels," in *Proc. USENIX Secur. Symp.*, 2006, Art. no. 5.
- [29] R. J. Walls, K. Kothari, and M. K. Wright, "Liquid: A detection-resistant covert timing channel based on IPD shaping," *Comput. Netw.*, vol. 55, no. 6, pp. 1217–1228, 2011.
- [30] S. Gianvecchio, H. Wang, D. Wijesekera, and S. Jajodia, "Model-based covert timing channels: Automated modeling and evasion," in *Proc. Int. Workshop Recent Adv. Intrusion Detection*, Springer, 2008, pp. 211–230.
- [31] F. V. Yarochkin, S. Dai, C. Lin, Y. Huang, and S. Kuo, "Towards adaptive covert communication system," in *Proc. 14th IEEE Pacific Rim Int. Symp. Dependable Comput.*, 2008, pp. 153–159.
- [32] J. Fridrich, *Steganography in Digital Media: Principles, Algorithms, and Applications*. Cambridge, U.K.: Cambridge Univ. Press, 2009.

- [33] M. Kharrazi, H. T. Sencar, and N. D. Memon, "Cover selection for steganographic embedding," in *Proc. Int. Conf. Image Process.*, 2006, pp. 117–120.
- [34] Z. Wang, G. Feng, L. Shen, and X. Zhang, "Cover selection for steganography using image similarity," *IEEE Trans. Dependable Secur. Comput.*, vol. 20, no. 3, pp. 2328–2340, May/June 2023.
- [35] O. Evsutin, A. Kokurina, and R. V. Meshcheryakov, "Approach to the selection of the best cover image for information embedding in JPEG images based on the principles of the optimality," *J. Decis. Syst.*, vol. 27, no. Supp, pp. 256–264, 2018.
- [36] H. Sajedi and M. Jamzad, "Cover selection steganography method based on similarity of image blocks," in *Proc. Int. Conf. Comput. Inf. Technol. Workshops*, 2008, pp. 379–384.
- [37] H. Sajedi and M. Jamzad, "Using contourlet transform and cover selection for secure steganography," *Int. J. Inf. Secur.*, vol. 9, no. 5, pp. 337–352, 2010.
- [38] S. Wu, Y. Liu, S. Zhong, and Y. Liu, "What makes the stego image undetectable?," in *Proc. Int. Conf. Internet Multimedia Comput. Serv.*, 2015, pp. 47:1–47:6.
- [39] Z. Zhou, Y. Mu, and Q. M. J. Wu, "Coverless image steganography using partial-duplicate image retrieval," *Soft Comput.*, vol. 23, no. 13, pp. 4927–4938, 2019.
- [40] H. Liu et al., "To deliver more information in coverless information hiding," *Multimedia Tools Appl.*, vol. 83, no. 3, pp. 7215–7229, 2024.
- [41] L. Ji, Y. Fan, and C. Ma, "Covert channel for local area network," in *Proc. IEEE Int. Conf. Wireless Commun. Netw. Inf. Secur.*, 2010, pp. 316–319.
- [42] T. Schmidbauer, S. Wendzel, A. Mileva, and W. Mazurczyk, "Introducing dead drops to network steganography using ARP-caches and SNMP-walks," in *Proc. 14th Int. Conf. Availability Rel. Secur.*, 2019, Art. no. 64.
- [43] W. Mazurczyk and K. Szczypiorski, "Steganography of VoIP streams," in *Proc. OTM Confederated Int. Conf. "On Move Meaningful Internet Systems"*, R. Meersman and Z. Tari, Eds., Springer, 2008, pp. 1001–1018.
- [44] L. Y. Bai, Y. Huang, G. Hou, and B. Xiao, "Covert channels based on jitter field of the RTP header," in *Proc. Int. Conf. Intell. Inf. Hiding Multimedia Signal Process.*, 2008, pp. 1388–1391.
- [45] L. Ying, Y. Huang, J. Yuan, and L. Y. Bai, "A novel covert timing channel based on RTP/RTCP," *Chin. J. Electron.*, vol. 21, no. 4, pp. 711–714, 2012.
- [46] S. Wendzel and J. Keller, "Hidden and under control - A survey and outlook on covert channel-internal control protocols," *Ann. des Télécommunications*, vol. 69, no. 7/8, pp. 417–430, 2014.
- [47] L. Caviglione, W. Mazurczyk, and S. Wendzel, "Advanced information hiding techniques for modern botnets," in *Botnets: Architectures, Countermeasures, and Challenges*, Boca Raton, FL, USA: CRC Press, 2019, pp. 165–188.
- [48] J. Lubacz, W. Mazurczyk, and K. Szczypiorski, "Hiding data in VoIP," in *Proc. 26th Army Sci. Conf.*, 2008.
- [49] V. Berk, A. Giani, and G. Cybenko, "Detection of covert channel encoding in network packet delays," Dartmouth College, Hanover, NH, Tech. Rep. TR2005-536, 2005.
- [50] X. Zhang, Y. Tan, C. Liang, Y. Li, and J. Li, "A covert channel over VoLTE via adjusting silence periods," *IEEE Access*, vol. 6, pp. 9292–9302, 2018.
- [51] T. Schmidbauer and S. Wendzel, "SoK: A survey of indirect network-level covert channels," in *Proc. ACM Asia Conf. Comput. Commun. Secur.*, 2022, pp. 546–560.
- [52] H. Schulzrinne, S. L. Casner, R. Frederick, and V. Jacobson, "RTP: A transport protocol for real-time applications," RFC 3550, Jul. 2003.
- [53] F. Lehner, W. Mazurczyk, J. Keller, and S. Wendzel, "Inter-protocol steganography for real-time services and its detection using traffic coloring approach," in *Proc. IEEE 42nd Conf. Local Comput. Netw.*, 2017, pp. 78–85.
- [54] A. Ovadya, R. Ogen, Y. Mallah, N. Gilboa, and Y. Oren, "Cross-router covert channels," in *Proc. 13th USENIX Workshop Offensive Technol.*, 2019.
- [55] G. C. Clarke Jr. and J. B. Cain, *Error-Correction Coding for Digital Communication*. Berlin, Germany: Springer Science Business Media, 1981.
- [56] L. R. Vermani, *Elements of Algebraic Coding Theory*. Evanston, IL, USA: Routledge, 1996.
- [57] K. Lamshöft and J. Dittmann, "Assessment of hidden channel attacks: Targetting Modbus/TCP," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 11100–11107, 2020.
- [58] F. J. Massey, "The Kolmogorov-Smirnov test for goodness of fit," *J. Amer. Statist. Assoc.*, vol. 46, no. 253, pp. 68–78, 1951.
- [59] W. Mazurczyk, S. Wendzel, S. Zander, A. Houmansadr, and K. Szczypiorski, *Information Hiding in Communication Networks: Fundamentals, Mechanisms, and Applications*. Hoboken, NJ, USA: Wiley, 2016.
- [60] S. Zillien, T. Schmidbauer, M. Kubek, J. Keller, and S. Wendzel, "Look what's there! Utilizing the Internet's existing data for censorship circumvention with OPPRESSION," in *Proc. 19th ACM ASIA Conf. Comput. Commun. Secur.*, 2024.



Steffen Wendzel (Member, IEEE) received the PhD and Habilitation degree from FernUniversität in Hagen, Germany, in 2013 and 2020, respectively. He is a professor with Hochschule Worms, Germany, where he is the scientific director of the Center for Technology and Transfer. He is also a lecturer with the Faculty of Mathematics & Computer Science, FernUniversität in Hagen. Before joining Hochschule Worms, he led a smart building security research team with Fraunhofer FKIE in Bonn, Germany. He (co-)authored more than 180 publications and (co-)

organized several conferences as well as special issues for major journals. His major research focus is on censorship circumvention, scientific taxonomies, and IoT security. Website: <https://www.wendzel.de>.



Tobias Schmidbauer received the diploma degree in computer science for public management from the University of Applied Science Hof, Germany, in 2016, and the MSc and PhD degrees from the Faculty of Mathematics & Computer Science, FernUniversität in Hagen, Germany, in 2019 and 2023, respectively. He is a professor with the Nuremberg Institute of Technology. Alongside his studies, he worked from 2007 until 2017 for the datacenter of the tax administration of the German federal state of Bavaria. Since 2017, he worked full-time for the

Bavarian State Office for Information Security until April 2024.



Sebastian Zillien received the BSc degree in applied informatics, in 2018, and the MSc degree in mobile computing, in 2020 from Hochschule Worms. He is currently working toward the PhD degree with the Hochschule Worms, Germany. He is a researcher with the Hochschule Worms, Germany, where he works for the projects WoDiCoF+ and ATTRIBUT. His research focus includes network and protocol-level security, anomaly detection, covert channels and reliability. He has had publications, among others, with AsiaCCS, *IEEE Transactions on Dependable and*

Secure Computing, IFIP SEC, NordSec and *Journal of Universal Computer Science*.



Jörg Keller received the MSc, PhD, and habilitation degrees from Saarland University, Saarbrücken, in 1989, 1992, and 1996, respectively. He is a professor of computer engineering with FernUniversität in Hagen, Germany, where he leads the Parallelism & VLSI research group since 1996. His research interests include network steganography, cryptographic primitives for embedded systems, energy-efficient and fault-tolerant parallel computation, and blended and virtual laboratories. He is author or co-author of 2 books and more than 180 refereed articles in journals

and conference proceedings. He has co-organized numerous conferences and workshops, and special issues. Website: <https://feu.de/pv/en>